The Optimal Structure of Conservation Agreements and Monitoring*

Heidi Gjertsen
San Diego Hunger Coalition

Theodore Groves
University of California, San Diego

David A. Miller
University of Michigan

Eduard Niesten
Conservation International

Dale Squires
NOAA/NMFS Southwest Fisheries Science Center

Joel Watson
University of California, San Diego

April 21, 2016

*The authors thank the NOAA Fisheries Service for supporting this research. The initial step of this research project was presented by Groves at the 15th Annual Conference of The African Econometric Society, 7–9 July 2010. Groves is especially grateful to Professor William Mikhail for his many courtesies and gracious hospitality. The authors also appreciate helpful comments from Bård Harstad, seminar participants at UC Berkeley, and conference participants at the 2016 University of Oslo Theory and Environment Workshop. Miller and Watson thank Isla Globus-Harris, Jake Johnson, and Erik Lillethun for assistance with analysis of the theoretical model.

Abstract

We examine the structure and performance of conservation agreements, which are used across the world to protect natural resources. Key elements of these agreements are: (1) they are ongoing arrangements between a local community and an outside party, typically a non-governmental organization (NGO); (2) they feature payments in exchange for conservation services; (3) the prospects for success depend on the NGO engaging in costly monitoring to detect whether the community is foregoing short-term gains to protect the resource; (4) lacking a strong external enforcement system, they rely on self-enforcement; and (5) the parties have the opportunity to renegotiate at any time. We provide a novel model that contains these ingredients and we apply the model to assess the workings of real conservation agreements, using three case studies as representative examples. We characterize equilibrium play (including how punishments and rewards are structured) and we show how the parties' relative bargaining powers affect their ability to sustain cooperation over time. The model captures important features of real conservation agreements and identifies some of the features required for successful agreements.

Contents

1 Introduction **2**

2 The Model **5**

2.1 Cooperation requires future punishments and rewards . . . 8
2.2 Structure of agreement play . . . 10
2.3 Structure of disagreement play . . . 12
2.4 Bargaining outcomes . . . 13
2.5 Summary of contractual equilibrium conditions . . . 14
2.6 Comparative statics and implications . . . 16
2.7 Comments on stock dynamics . . . 18

3 Interpretation and Implications **20**

4 Case Studies **25**

4.1 Forest protection in Cambodia . . . 25
4.2 Laos deer conservation . . . 29
4.3 Grey whale habitat protection in Mexico . . . 31

5 Conclusion **34**

A Appendix **35**

1 Introduction

Incentive-based programs are increasingly employed by conservation practitioners to encourage changes in the use of environmental resources.[1] While past approaches generally focused on negative incentives such as fines and penalties, many current programs use a positive incentive, namely compensation of various forms to encourage conservation activities. These approaches recognize that conservation can entail a loss in terms of foregone income or access to resources.

Conservation agreements constitute an important class of positive incentive systems. Under this approach, conservation investors (typically non-governmental organizations, NGOs) negotiate quid-pro-quo contracts by which resource users forego destructive activities in exchange for benefits provided by the investors. Monitoring of the resource users is required, so that benefits are conditioned on some measure of performance. Conservation agreements involve ongoing interaction between the resource users and conservation investors in the absence of a strong external enforcement institution. Thus, they rely on self-enforcement, whereby the investors provide benefits periodically over time conditional on verification that conservation performance targets are met. Benefits may be in the form of cash, services, or goods.

The nature of this relationship is fairly unique among conservation interventions. Other interventions are either one-shot agreements or they do not link payments to successful conservation. For instance, the typical community-based project involves an external flow of funds to the community (e.g. through a grant), but the funds are not explicitly contingent on whether the community successfully protects a specified natural resource. Outright land purchases and timber concessions share the characteristic of compensating for opportunity costs, but they do not involve a continuing relationship between conservation investors and resource users.

Conservation agreements have been increasingly adopted worldwide. For example, Conservation International's Conservation Stewards Program has systematically implemented 51 such programs in 14 countries. Other international conservation organizations, such as the Wildlife Conservation Society and The Nature Conservancy, have experimented with the approach, as have many smaller local conservation organizations (TNC (2013); Niesten and Gjertsen (2010); Svadlenak-Gomez, Clements, Foley, Kazakov, Lewis, Miguelle, and R.Stenhouse (2007)). Given that experimentation is fairly recent, there is little empirical evidence to suggest how they perform or how they are best designed and implemented,

[1]The following publications document the range of incentive programs: Ferraro (2001); Ferraro and Kiss (2002); Milne and Niesten (2009); Simpson and Sedjo (1996); Troëng and Drews (2004); Wunder (2004, 2008).

though conservation practitioners have begun to address these questions (Niesten, Bruner, Rice, and Zurita 2008; TNC & CI 2009). Also, little theoretical literature has been developed on this topic. Exceptions include analysis of moral hazard in conservation contract design (Ferraro a; Hart, Latacz, and Lohmann; Latacz, Lohmann, and der Hamsvoort; Wu and Babcock), cost-effectiveness of conservation payments (Ferraro b), marine conservation easements (Deacon and Parker), and general conceptual models of payments for environmental services (Engel and Palmer; Engel, Pagiola, and Wunder). However, none of these studies directly address the repeated nature of the interactions between the parties to the contract.

This paper aims to explore the incentive problems for conservation agreements by developing a theoretical model and then using the model to evaluate how agreements have worked in three case studies. Key aspects of real conservation agreements are that they (i) entail ongoing relationships between local communities and NGOs, (ii) lack strong external enforcement and thus function mainly on the basis of self-enforcement, (iii) require costly monitoring to detect whether the communities are fulfilling their promises to protect environmental resources, and (iv) can be renegotiated by the parties at any time.

To capture these features, our model specifies a repeated game between a Community and an NGO, with an explicit account of bargaining and transfers within each period. The model has several novel aspects. First, we incorporate equilibrium selection and bargaining power using the *contractual equilibrium* solution concept (Miller and Watson 2013; Watson 2013). In a contractual equilibrium, an endogenous disagreement point is identified for each public history of play, and the parties share the surplus according to fixed bargaining weights that represent the details of the bargaining protocol. Second, in the stage game, while the Community chooses whether to protect the resource, costly monitoring by the NGO is required to provide a public signal of the Community's choice. Thus, there is an informational asymmetry and incentive problems on both sides. Third, the model allows for the parties to have different discount factors.

We characterize equilibrium play, including how punishments and rewards are structured, and we examine the relation between the joint value attained in equilibrium and the parameters of the relationships, such as the parties' relative bargaining power, the cost of monitoring, the benefits of preserving or exploiting the natural resource, and the discount factors. To explore how the ingredients of contractual equilibrium translate into the real world, we offer a number of general implications. We also discuss how the contractual equilibrium may be interpreted as a series of short-term agreements linked by the parties' expectations over time, which, as our case studies illustrate, corresponds to the manner in which real conservation agreements are managed.

Our results regarding bargaining power may be of particular interest. We find that it is generally optimal for the parties to specify punishment paths that are less efficient than is their desired cooperation path, but they anticipate renegotiating to achieve the joint value of cooperation. In equilibrium, the effective punishments (incorporating renegotiation) must be sufficient to provide the parties with the proper incentives to cooperate. Because the effective punishments depend on how the parties share in the surplus of renegotiation, bargaining power plays an important role in determining the severity of the punishments and the achievable level of cooperation. We show, in particular, that the joint value is increasing in the community's bargaining power. The higher joint value is associated with a lower intensity of monitoring by the NGO.

On the theoretical side, our modeling exercise contributes to the growing literature on various self-enforced environmental agreements using repeated game models (examples include **????**). Much of this literature focuses on familiar ideas from repeated game theory, including the folk theorem. The papers that consider negotiation (including the four just noted) use the abstract notion of "renegotiation proofness" (Bernheim and Ray 1989; Farrell and Maskin 1989), for which bargaining power plays no role and negotiation is not modeled directly. Thus, our analysis of bargaining power provides a new element for the theory of environmental agreements and for the evaluation of actual agreements. Also, our model highlights the importance of providing incentives to monitor compliance, and it handles cases in which the parties have different discount factors. Further, we offer a modest extension to demonstrate how the dynamics of the resource stock may interact with incentives, although a full-blown analysis of stock dynamics is beyond the scope of this paper.

In our model, the NGO's choice of whether to monitor can be regarded as a technology choice. **?** analyze a different type of technology choice within a period—one that affects the costs and benefits of emissions selected later in the period.[2] In another related vein, **?** looks at an ongoing conservation choice by the owner of a natural resource who can sell or lease it to a prospective buyer, finding conditions under which a lease arrangement is preferred. This may explain the prolific nature of conservation agreements and can be seen as motivation for our study herein.[3]

Our contribution on the applied side is to use the implications of the model to evaluate actual conservation agreements. We offer three case studies that differ in terms of success,

[2]**?** analyze how the technology choices interact with emission choices in subgame perfect equilibria. For moderate discount values, technology choices may be higher or lower than is efficient. **?** examine how a long-term technology choice affects incentives to cooperate in a relational contract.

[3]In Harstad's model, the prospective buyer values the existence of the non-depleted resource. The buyer can ensure ongoing conservation by purchasing the resource, but only if the resource is conserved prior to the purchase. The seller is willing to conserve in order to sell the resource to the buyer, but if the seller would conserve over time then the buyer has no need to purchase. Inefficiency is inescapable in equilibrium.

or lack thereof, and we discuss possible reasons for their outcomes. The first case, which we rate as a success, involved protection of a forest in Cambodia; the second (an unsuccessful attempt) dealt with an endangered species of deer in Laos; and the third (an ongoing success) seeks to preserve the marine habitat of whales in Baja California, Mexico. Characteristics of each case are compared with the ideal conditions for cooperation and the contractual specifications that our model suggests.

The next section defines the repeated game model and derives the main theorems characterizing an equilibrium. Section 3 interprets the model as it applies to conservation agreements and summarizes the characteristics of optimal agreements. Section 4 contains the real-world case studies. We conclude in Section 5 with some remarks about future work and applications. Appendix A provides a detailed construction of the equilibrium described in Section 2.

2 The Model

A Community (C) and an NGO (N) interact in discrete periods of time over an infinite horizon. In each period, there are two phases:

- the *bargaining phase*, where the players negotiate on how to coordinate their future behavior and can also make immediate monetary transfers; and
- the *action phase*, where productive interaction occurs.

In the action phase, the players interact according to the following stage game:

		NGO	
		M	R
Community	P	$0, b-c$	$0, b$
	E	$e, -c$	$e, 0$

The Community chooses to either "protect" its local natural resource (P) or "exploit" it (E). If the Community exploits, it obtains a gain of $e > 0$ in the period. If the Community protects, then it obtains no gain but the NGO (on behalf of its donors) earns a benefit of $b > 0$. Simultaneously, the NGO can either "monitor" the Community (M) or "rest" (R). The cost of monitoring is $c > 0$. The four action profiles are thus PM, PR, EM, and ER.

The informational assumptions in our model are designed to capture an important feature of many conservation agreements: that some sort of monitoring is required to observe whether the Community is taking the desired action to conserve the natural resource. Importantly, the NGO (and society) cares about the Community's action whether or not it is observed immediately—that is, the NGO's payoff depends on the Community's action either way—but to detect exploitation in the current period, the NGO must engage in monitoring. For example, suppose the Community is a village that chooses whether to protect a turtle nesting site. The Community's action influences the long-term viability of the turtle population, which the NGO cares about. In the long run, the NGO may be able to estimate the Community's actions over time by evaluating the health of the turtle population, but this assessment entails a great deal of noise and time lag. If the NGO wants a signal of the Community's action within a period, it will have to send a worker to the village to observe and record the Community's action.[4]

Thus, we assume that the NGO's action is publicly observed but the Community's action is private. By monitoring, the NGO obtains information about the Community's action. Specifically, at the end of the period there is a monitoring signal $s \in \{G, B\}$, where G stands for "good" and B stands for "bad." The bad signal means that the NGO finds evidence that the Community chose to exploit the natural resources. If the Community selects P then the signal is $s = G$ for sure. If the NGO selects R then $s = G$ for sure as well; that is, the NGO cannot detect exploitation without monitoring. However, if the NGO selects M and the Community selects E, then $s = B$ with probability $\lambda > 0$ and $s = G$ with probability $1 - \lambda$. In the case of $\lambda < 1$, there is some noise in the monitoring technology, so exploitation is not always detected. Assume that both the NGO and the Community observe the signal, so at the end of a period the possible publicly observed outcomes are RG, MG, and MB.[5]

Because the parties can make transfers, welfare is given by the sum of their payoffs—their *joint value*. We assume that $b - c > e$, so that it is more efficient to have the action profile PM, where the Community protects and the NGO monitors, than to have action profile ER, where neither occurs. The most efficient action profile is clearly PR. That is, monitoring does not create a direct benefit to either party. Rather, monitoring generates information that randomly identifies an incident of exploitation. The structure of payoffs in this game is shown in Figure 1.

Our solution concept is a refinement of perfect public equilibrium called *contractual equilibrium* (Miller and Watson 2013; Watson 2013), where behavior in the action phase is

[4]Note that it is an assumption of equilibrium that the NGO knows what actions the Community plans to take, so the NGO can evaluate its preferences over different equilibria without observing the Community's actions.

[5]Our results would hold with a more general monitoring choice and signal, although a greater range of outcomes would occur with positive probability on the equilibrium path (for instance, if $s = B$ would be realized with positive probability even when the Community chooses P).

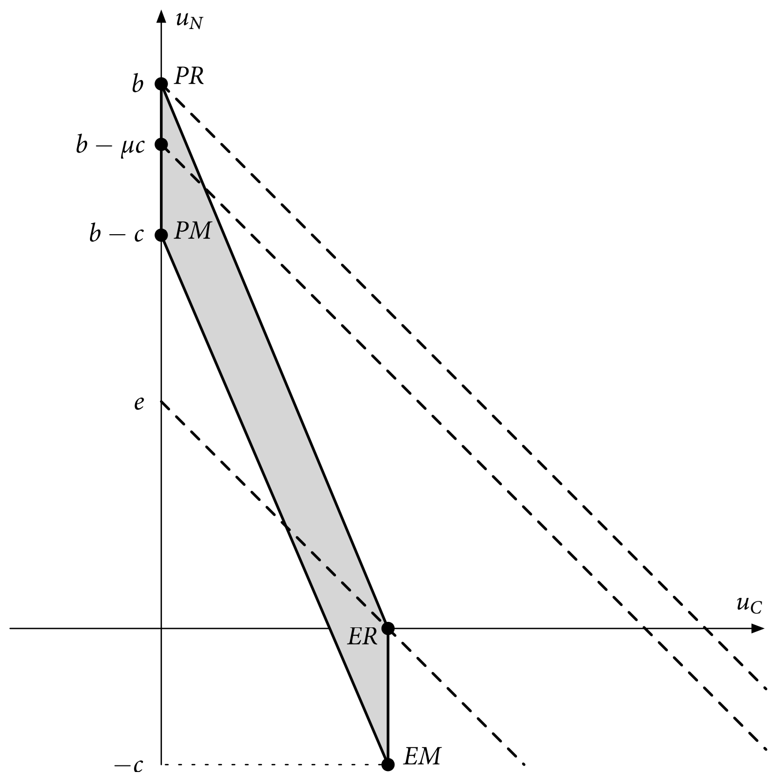

FIGURE 1. Stage game payoffs. The grey region is attainable without monetary transfers, using mixtures of action profiles PR, ER, PM, and EM. Heavy dashed lines illustrate payoff vectors that are attainable with monetary transfers combined with pure action profiles PR, ER, and a mixture of PR with probability μ and PM with probability $1 - \mu$, for a given $\mu \in (0, 1)$.

consistent with individual incentives and the outcome of bargaining each period is given by a standard bargaining solution with exogenous bargaining weights.[6] Given the absence of an effective legal system, there is no external enforcement, so any agreement must be self-enforced. That is, the players can sustain cooperation only by appropriately rewarding or punishing each other over time.

To consider incentives within a period, it is convenient to write each player's value from the beginning of a given period as the sum of its payoff within the period and its discounted continuation value from the start of the next period. Let $\delta_C < 1$ and $\delta_N < 1$ denote the players' discount factors, and let v_C^t and v_N^t denote their continuation values from the negotiation phase of a given period t. Suppose that in period t the NGO makes a monetary transfer of m^t to the Community. Let u_C^t and u_N^t denote their stage-game payoffs in period t, and let v_C^{t+1} and v_N^{t+1} denote their continuation values from the start of period $t+1$, in total discounted terms. We thus have

$$v_C^t = m^t + u_C^t + \delta_C \mathbb{E}\big(v_C^{t+1}\big)$$

and

$$v_N^t = -m^t + u_N^t + \delta_N \mathbb{E}\big(v_N^{t+1}\big).$$

Note that we treat v_C^{t+1} and v_N^{t+1} as random variables because they can be conditioned on the public outcome of interaction in period t. We can begin to analyze incentives in period t by considering how this conditioning can be structured.

2.1 Cooperation requires future punishments and rewards

First consider a scenario in which in period t the players regard v_C^{t+1} and v_N^{t+1} as fixed and independent of the public outcome in period t. For instance, this would be the case if when the players negotiate in period t they discuss plans only for period t but not for the future, and they expect to start "from scratch" again in the next period. Then the Community will not be motivated to select P, because deviating to E would raise its stage-game payoff from $u_C^t = 0$ to $u_C^t = e$. Even if the NGO made an up-front payment in exchange for protecting the natural resource, the payment is sunk before the action phase and so it would not affect the Community's incentives in the stage game.

[6] Contractual equilibrium represents an explicit account of bargaining within each period and incorporates a theory of disagreement. The construction appears later in this section. The previous literature focused on the more abstract notion of "renegotiation proofness" (Bernheim and Ray 1989; Farrell and Maskin 1989), in which bargaining power plays no role. We characterize contractual equilibrium using a recursive formulation of continuation values, following **?**.

We conclude that motivating the Community to protect the natural resource requires the parties to condition their future behavior on the outcome of the action phase in period t. In this way, the continuation value v_C^{t+1} will depend on whether the Community chooses E or P in period t, so that the Community is rewarded for choosing P and punished for choosing E. The reward is the promise of future transfers, which raise the continuation value v_C^{t+1}. It is important to recognize that it is not the transfer m^t that motivates the Community to choose P in period t; rather, it is the prospect of future transfers.

But there is an additional problem: The Community's continuation payoff cannot depend directly on its action in the stage game, because its action is not publicly observed. Rewarding and punishing the Community requires the NGO to engage in monitoring, and monitoring entails an additional incentive condition. Because monitoring is costly, the NGO must be rewarded in future periods for monitoring in period t; that is, the continuation value v_N^{t+1} must depend on the NGO's action in period t. Further, since monitoring is noisy, the rewards and punishments for both the Community and NGO are subject to random noise and thus are figured as expectations.

Thus, there are incentive problems on both sides—on the Community's side with respect to protecting the natural resource, and on the NGO's side with respect to whether to monitor. The parties would like to solve these incentive problems with as little monitoring as possible.

In a contractual equilibrium, the players always anticipate that they will achieve their highest attainable joint value from the start of the next period, regardless of the state of affairs in the current period. This is because they are free to negotiate at the beginning of each period and they can make arbitrary transfers. If their default specification of behavior from the start of the next period would yield a lower joint value, then the players would expect to renegotiate up to the attainable frontier.

To describe the set of possible continuation values, let L denote the highest attainable expected joint value. As just argued, in a contractual equilibrium the continuation values always satisfy $v_C^{t+1} + v_N^{t+1} = L$. Thus, we know that the set of possible continuation values has the form

$$V = \{\sigma z^C + (1 - \sigma) z^N \mid \sigma \in [0, 1]\},$$

where vectors $z^C = (z_C^C, z_N^C)$ and $z^N = (z_C^N, z_N^N)$ have the properties that

$$z_C^C + z_N^C = L \tag{1}$$

$$z_C^N + z_N^N = L \tag{2}$$

and $z_N^C \geq z_N^N$.

This means that the set of possible continuation value vectors (again, in total discounted terms) is a line segment with slope -1, such that z^C is the point that most favors the NGO (punishes the Community) and z^N is the point that most favors the Community.[7] Different points on this line segment correspond to different transfers between the players in period $t+1$ and beyond. Define the *span* of V to be the difference d between the Community's payoffs at the endpoints of V:

$$d = z^N_C - z^C_C. \tag{3}$$

Then $d = z^N_N - z^C_N$ as well by Eq. 1–2, and we can write $V = \{z^C + (\zeta, -\zeta) \mid \zeta \in [0, d]\}$.

2.2 Structure of agreement play

We can think of an agreement between the parties in period t as specifying the immediate transfer m^t, an action profile to be played in period t, and a function that relates the continuation values $(v^{t+1}_C, v^{t+1}_N) \in V$ to the public outcome of period t. The continuation values represent how the parties will coordinate in the future, including the future monetary transfers that the NGO will make.

The best agreement for the players—that is, the one that maximizes their joint value in equilibrium—has the Community protecting the natural resource (selecting P) and the NGO randomizing over whether to monitor (choose M); the NGO selects M with the probability $\mu \in (0, 1)$ that is calculated below. If the public outcome is RG, which means the NGO did not monitor, then the parties coordinate on the continuation value vector $z^N = z^C + (d, -d)$, which is worst for the NGO in period $t+1$. If the public outcome is MB, which means the Community deviated, then the parties coordinate on the continuation value vector z^C, which is worst for the Community in period $t+1$. If the public outcome of the stage game is MG, so that the NGO monitored and there is no evidence of exploitation, then the parties coordinate on an intermediate continuation value vector $z^C + (x, -x)$ in period $t+1$, where $x \in [0, d]$ is calculated below.

For the NGO to monitor with probability $\mu \in (0, 1)$, it must be indifferent between monitoring and resting. Given that the Community selects P, the NGO's expected continuation value from the bargaining phase in period t when it monitors in this period is:

$$w_N = -m^t + b - c + \delta_N(z^C_N - x).$$

[7]The set V is convex because we assume that the players can use a public randomization device to achieve any convex combination of z^C and z^N.

The NGO's expected value of resting is

$$-m^t + b + \delta_N(z_N^C - d).$$

These must be equal, so we must have

$$x = d - \frac{c}{\delta_N}. \tag{4}$$

For the Community to be motivated to protect the natural resource, its expected value of selecting P must be greater than or equal to its value of selecting E in period t. The Community's value from the bargaining phase in period t when it selects P in this period is:

$$w_C = m^t + \delta_C \left[\mu(z_C^C + x) + (1-\mu)(z_C^C + d)\right].$$

The term in brackets is the expected continuation value from period $t+1$, given that the NGO randomizes with probability μ in period t and the outcome affects how the players coordinate in period $t+1$. The Community's expected value of choosing E is

$$m^t + e + \delta_C \left[\mu\lambda z_C^C + \mu(1-\lambda)(z_C^C + x) + (1-\mu)(z_C^C + d)\right].$$

In the bracketed part, the first term is the probability that the Community is caught exploiting the natural resource times its punishment continuation value in period $t+1$. The second term accounts for the chance that the NGO monitors but receives the good signal G, and the third term accounts for the chance that the NGO does not monitor.

For the Community's continuation value of P to exceed that of E, we need

$$\mu \geq \frac{e}{\delta_C \lambda x}.$$

Because monitoring is costly (and lowers the joint value), it is optimal to set μ as low as possible, which means

$$\mu = \frac{e}{\delta_C \lambda x}. \tag{5}$$

Since the Community and NGO always jointly obtain L in the agreement, it must be that $L = w_C + w_N$, so substituting in the expressions for w_C and w_N from above yields

$$L = b - c + \delta_N(z_N^C - x) + \delta_C\left[\mu(z_C^C + x) + (1-\mu)(z_C^C + d)\right]. \tag{6}$$

The players divide the joint value as desired by selecting an appropriate immediate transfer m^t. Their relative shares are determined by bargaining, as specified below.

2.3 Structure of disagreement play

Although the players always anticipate being able to negotiate to the attainable frontier of continuation values, an agreement in one period must describe how the players will behave in future periods in case they should fail to agree when renegotiating in the future. That is, the agreement in a given period specifies the players' disagreement points for later periods. Our theory of disagreement (from Miller and Watson 2013) is that when they disagree in period t, no transfer is made and they coordinate in some incentive-compatible way until the next agreement is made—which the players anticipate will occur in the following period. Because there are typically multiple incentive-compatible ways to coordinate, various disagreement points are possible.

The key to characterizing contractual equilibrium is to identify the disagreement point that most favors the Community and the disagreement point that most favors the NGO. The former will be useful in punishing the NGO and rewarding the Community, whereas the latter will be used to punish the Community and reward the NGO. In both cases, a disagreement point for period t is a payoff vector constructed from actions in the stage game at period t and a specification of the continuation value from period $t+1$ as a function of the public outcome of period t.

As shown in Appendix A, the disagreement point that most favors the Community involves playing ER in period t, followed by continuation value z^N from the start of period $t+1$, regardless of the actual outcome in period t. That is, in situations in which the NGO is to be punished (and the Community rewarded), if they should disagree then the Community exploits the natural resource and the NGO does not monitor. This specification yields continuation values $y^N = (y_C^N, y_N^N)$, where

$$y_C^N = e + \delta_C z_C^N \tag{7}$$

and

$$y_N^N = \delta_N z_N^N. \tag{8}$$

Appendix A also shows that the disagreement point that most favors the NGO involves playing PM in period t, followed by continuation value

$$z^C + \left(\frac{e}{\lambda\delta_C}, -\frac{e}{\lambda\delta_C}\right)$$

from the start of period $t+1$.[8] That is, in situations in which the Community is to be punished (and the NGO rewarded), if they should disagree then the Community is expected to protect the natural resource and the NGO monitors with probability 1. This specification yields continuation value $y^C = (y^C_C, y^C_N)$, where

$$y^C_C = \delta_C z^C_C + \frac{e}{\lambda} \tag{9}$$

and

$$y^C_N = b - c + \delta_N z^C_N - \frac{\delta_N e}{\delta_C \lambda}. \tag{10}$$

2.4 Bargaining outcomes

The final step in describing contractual equilibrium is to account for the outcome of bargaining at the beginning of each period. The players will always obtain the joint value L under agreement, but the division of this value will depend on the disagreement point, which itself may be determined by behavior in previous periods.

Our bargaining theory is that the players divide the bargaining surplus according to fixed bargaining weights $\pi = (\pi_C, \pi_N)$, as in the Nash (1950) bargaining solution with transfers. These weights represent the exogenously specified bargaining protocol for negotiation in each period. Thus, if $y = (y_C, y_N)$ is the vector of disagreement values, then the Community obtains $y_C + \pi_C(L - y_C - y_N)$ and the NGO obtains $y_N + \pi_N(L - y_C - y_N)$.

We can now establish a relation between points z^C and z^N and the disagreement values y^C and y^N. Recall that z^C is the Community's least favorite vector of continuation values from the start of a period. It must be supported by a Nash bargaining outcome relative to some disagreement point. Because the disagreement point that least favors the Community is y^C, we have

$$z^C = y^C + \pi\left(L - y^C_C - y^C_N\right). \tag{11}$$

Note that z^C, y^C, and π are vectors, so Eq. 11 is actually two scalar equations, one for each

[8] The additional amount $e/\lambda\delta_C$ in the Community's continuation value compensates the Community for protecting the natural resource in period t. If the Community deviates (exploits) and is caught—so the signal is B—then the parties coordinate on continuation value z^C from the start of period $t+1$. The Community is indifferent between protecting and exploiting, and thus is willing to protect.

player. The same construction holds for the NGO's least favorite continuation value:

$$z^N = y^N + \pi \left(L - y_C^N - y_N^N \right). \tag{12}$$

2.5 Summary of contractual equilibrium conditions

Equations 1–12 characterize the set of contractual equilibria.[9] The unique solution to these equations (for the case in which the players are sufficiently patient) is displayed graphically in Figure 2. Each contractual equilibrium specifies agreement and disagreement behavior, as a function of past behavior. Distinct contractual equilibria differ only in the specification of behavior in contingencies in which the players have never before reached an agreement. These differences determine only the initial division of the joint value—that is, the location of the first period's continuation value on the line segment between z^C and z^N. If, for instance, we assume that the default behavior in the absence of any prior agreement is to play the stage-game Nash equilibrium ER repeatedly, then their first agreement is to start at z^N.

Theorem 1. *Equations 1–12 have a unique solution. If the solution satisfies*

$$d \geq \frac{c}{\delta_N} + \frac{e}{\delta_C \lambda} \tag{13}$$

then the solution identifies the contractual equilibrium values $z^{C*}, z^{N*}, L^*, d^*, x^*,$ *and* μ^**. The the NGO's equilibrium payment to the Community in period* $t+1$ *depends on the outcome of the stage game in period* t *and is increasing in the following order: RG, MG, MB.*

If the solution to Equations 1–12 violates Condition 13, then the unique contractual equilibrium features repeated play of the stage-game Nash equilibrium ER*.*

To understand Condition 13, note that in the contractual equilibrium we have described, $d^* = \frac{c}{\delta_N} + \frac{e}{\delta_C \lambda \mu^*}$, and μ^* can be no more than 1. That is, to achieve cooperation in a period, the span must be sufficiently large to reward both the Community for protecting (not taking the exploitation benefit e in the current period) and the NGO for monitoring (not avoiding the cost c).

[9] Note that Equations 1–12 give a system of 12 scalar equations, since one of the two scalar equations in (11) is redundant with (1), and one of the two scalar equations in (12) is redundant with (2). There are also 12 variables: $z_C^C, z_N^C, z_C^N, z_N^N, L, d, \mu, x, y_C^C, y_N^C, y_C^N, y_N^N$.

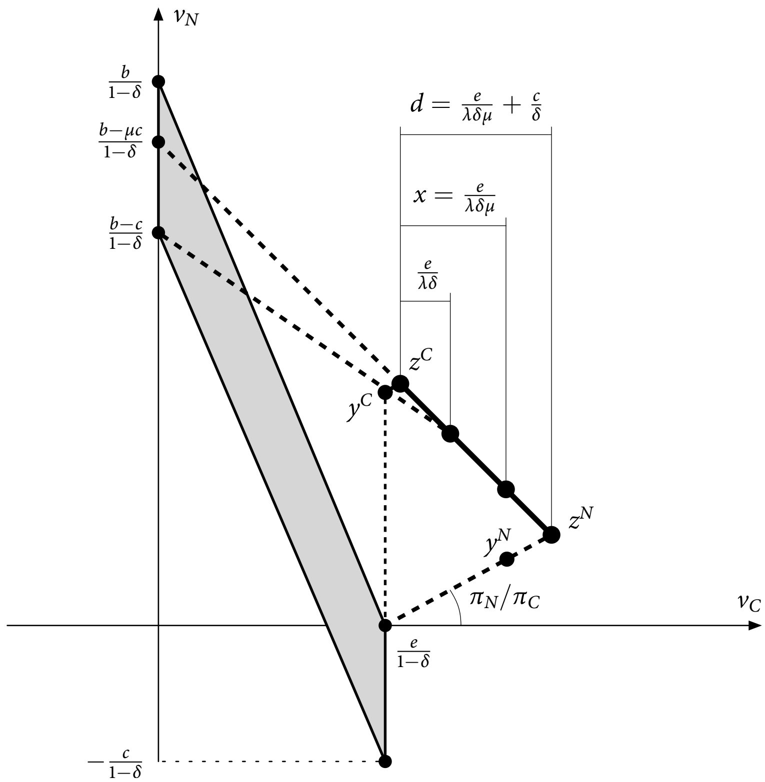

FIGURE 2. Contractual equilibrium for the case $\delta_N = \delta_C = \delta$. All payoffs are shown in total discounted terms. The contractual equilibrium value set V is attained along the equilibrium path. Its endpoints, z^C and z^N, are the expected payoffs in the states C and N, respectively. The payoff vectors y^C and y^N are attained under disagreement in the two states.

2.6 Comparative statics and implications

In this subsection, we provide results on how the contractual equilibrium values vary with the parameters of the model. The first of these results summarizes what we can say in general, without restricting the parameters.

Theorem 2. *Assume that Condition 13 for a non-degnerate contractual equilibrium holds. Then the contractual-equilibrium span d^* is increasing in b and λ, and decreasing in e and c. If also the players are sufficiently patient and the monitoring is sufficiently precise (δ_N, δ_C, and λ are sufficiently high), then in addition d^* is increasing in δ_N and δ_C.*

If the Community and the NGO have the same discount factor, then the equilibrium values take a simple form, and additional conclusions arise.

Theorem 3. *Assume that $\delta_C = \delta_N = \delta$ and Condition 13 for a non-degenerate contractual equilibrium holds. Then the equilibrium span is*

$$d^* = \frac{\pi_N e - \frac{e}{\lambda} + \pi_C(b-c)}{1-\delta},$$

the equilibrium monitoring probability is

$$\mu^* = \frac{e}{\lambda(\delta d^* - c)},$$

and the welfare level is

$$L^* = \frac{b - \mu^* c}{1-\delta}.$$

The span d^*, the probability of no monitoring $1-\mu^*$, and the welfare level are all increasing in π_C, λ, b, and δ; and decreasing in e and c.

The comparative statics with respect to π_C and $\pi_N = 1 - \pi_C$ are illustrated in Figure 2. The darker endpoints and contractual equilibrium value set correspond to a higher value of π_C, whereas the lighter endpoints and value set arise with a lower value of π_C. As the Community's bargaining weight increases, so does the joint value that the parties obtain. An increase in the Community's bargaining weight shifts the bargaining outcome in the Community's favor, but more so when the disagreement point is y^N than when it is y^C. The disagreement point y^N in the continuation most favorable to the Community (which punishes the NGO) is further from the frontier of the bargaining set than is the disagreement point y^C in the continuation favoring the NGO. Thus, in terms of enlarging the span needed to reward and punish the players, changes in relative bargaining power have a greater influence on the endpoint most favoring the Community. Therefore the parties jointly

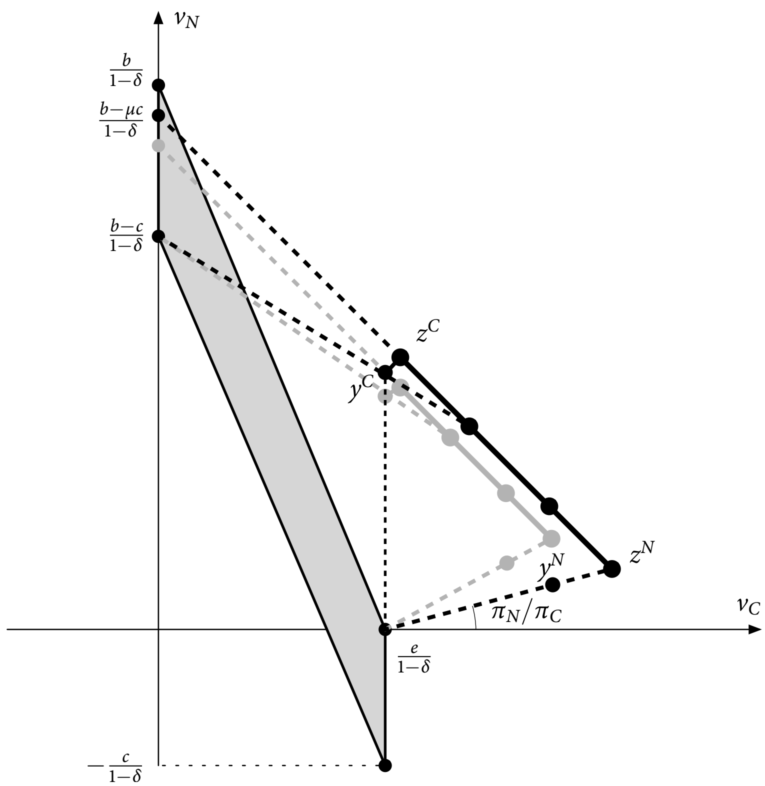

FIGURE 3. Comparative statics with respect to π_N/π_C, for the case $\delta_N = \delta_C = \delta$. As π_N/π_C decreases, the contractual equilibrium moves from the grey construction to the black construction. Because y^N is farther below the equilibrium joint value than is y^C, the increase in the Community's bargaining power has a greater effect on z^N than on z^C. As a consequence the span increases, providing more incentive power, so that less monitoring is needed on the equilibrium path.

prefer the Community's bargaining weight to increase, in order to increase the span of the contractual equilibrium value set.

This analysis pertains only to the parties' joint value. If at the beginning of the entire game, the parties negotiate subject to an exogenously fixed disagreement point (such as the "business as usual" Nash equilibrium of the stage game), then the NGO's shared interest in the Community having bargaining power is tempered by the fact that the initial share of surplus is sensitive to the parties' relative bargaining weights. To be more precise, as a thought exercise let us start by imagining $\pi_C = 0$ and $\pi_N = 1$. In this case, cooperation is not possible and so the parties are stuck with repeated selection of the stage-game Nash profile ER.[10] If we imagine raising π_C from 0 (and correspondingly lowering π_N from 1), then the attainable joint value increases; even though the NGO's share of surplus goes down, both the NGO and the Community are better off. But when π_C becomes large, although the joint value continues to rise, the NGO's selected equilibrium payoff from the beginning of the game eventually starts to decrease. In other words, the NGO likes the idea of giving the Community some bargaining power, but only up to a point, whereas society prefers that the Community's bargaining power be as large as possible.[11]

2.7 Comments on stock dynamics

For simplicity our model assumes a stationary setting, where the parameters of the stage game are unchanged over time and are thus not affected by past behavior. The model is best suited for settings in which the stock of the natural resource would not drastically change from period to period. In many real settings, natural resources follow a dynamic process and can be depleted or even crash, so it is useful (i) to determine whether our modeling approach can be extended to settings with a stock variable and (ii) to consider the possible implications of stock dynamics on incentives and contractual-equilibrium outcomes. While a general analysis entails complications that must be left for future papers, here we can provide some reassurance regarding how our techniques extend, and we can also work through some intuition.

On the technical front, the equilibrium construction can be generalized by writing the set of contractual-equilibrium continuation values V as a function of the environmental

[10] To see this graphically, take $\pi_N/\pi_C \to \infty$ in Figure 3. As the line through $\left(\frac{e}{1-\delta}, 0\right)$ and z_N rotates counterclockwise, the span of the contraction-equilibrium value set decreases and eventually collapses to zero, in which case there is no scope for continuation values to vary. To see this mathematically, substitute $\pi_C = 0$ and Eq. 8 into Eq. 11, which yields $z_C^N = e + \delta_C z_C^N$, or $z_C^N = \frac{e}{1-\delta_C}$. Thus the endpoint z^N that is supposed to favor the Community yields the same utility as repeated play of the stage game Nash equilibrium ER, leaving no scope for incentives.

[11] Better yet, the NGO would like to exercise a lot of bargaining power when first negotiating an agreement with the Community, but be able to commit to a low bargaining weight for all future negotiations.

resource stock, which we call the *state* and denote by θ. Because the players can negotiate and make transfers at the beginning of each period, $V(\theta)$ does not depend on the history except through θ; further, it is a line segment with slope -1. The key difference between our main model and the stock-variable generalization is that, in the latter, $V(\theta)$ will generally be different than $V(\theta')$ for $\theta \neq \theta'$. When expressing how continuation values from the start of the current period relate to the selection of continuation values from the next period (in both agreement and disagreement play), the analysis is similar to what is described above in this section. However, continuation values from the next period may be selected from multiple $V(\theta)$ sets because different action profiles in the current period may lead to different stock values in the following period. It is straightforward to define contractual equilibrium, but technical assumptions would be necessary to ensure existence.[12]

On the applied front, stock dynamics could help or hinder cooperation depending on how stage-game actions influence stock adjustments. One particular special case can be analyzed easily and illustrates how the possibility of stock collapse can enhance incentives. Consider the following variation of our model, which we call the *simple stock extension*. Assume that the parties have the same discount factor δ. Suppose that the resource stock remains healthy over time (from each period to the next) as long as the Community protects the resource. However, if the Community exploits the resource in a given period, then the resource stock recovers to a healthy state with probability β and permanently crashes with probability $1 - \beta$. The parties jointly observe whether a crash occurs. Assume that a crash would render the relationship worthless to both parties, so their continuation values would then be zero. That is, there are two states: $\overline{\theta}$ denoting a healthy stock and $\underline{\theta}$ denoting a stock that has crashed. We have $V(\underline{\theta}) = \{(0, 0)\}$.

In this stylized extension of our model, the contractual equilibrium is characterized just as in the main model, but with two modifications. First, the characterization of the disagreement continuation value y^N must be modified to incorporate the possibility of a stock crash following the Community's choice to exploit. Specifically, Equations 7 and 8 become

$$y_C^N = e + \delta\beta z_C^N$$

and

$$y_N^N = \delta\beta z_N^N,$$

because with probability $1-\beta$ the stock crashes, leading to continuation values of zero. Sec-

[12]Both the levels and spans of state-contingent sets of continuation values are interdependent. Suitable monotonicity conditions must be achieved to ensure a fixed point. Such conditions are well beyond the scope of this paper and have not yet been addressed in the abstract. We think this is a good topic for pure and applied theoretical work in the future.

ond, regarding agreement play, the Community's value of deviating to exploit the resource (that is, choosing E) becomes

$$m^t + e + \delta\beta\left[\mu\lambda z_C^C + \mu(1-\lambda)(z_C^C + x) + (1-\mu)(z_C^C + d)\right].$$

This alters the indifference condition that identifies μ, and Equation 5 becomes

$$\mu = \frac{e - \delta(1-\beta)(z_C^C + d)}{\delta(1-\beta)(x-d) + \delta\beta\lambda x}.$$

Lowering the recovery parameter β (that is, raising the probability that exploitation leads to a stock crash) has two opposing effects. First, it loosens the Community's incentive condition by making the choice of E less attractive. This effect would lower the equilibrium monitoring probability and contribute to a higher welfare level. Second, lowering β has a direct negative effect on y_C^N, which reduces the equilibrium span and makes it more difficult to reward and punish the players. The latter effect is weak if π_C is relatively large. Thus, we obtain the following result. See Appendix A for details.

Theorem 4. *In the simple stock extension, for π_C sufficiently large, reducing the recovery parameter β causes the contractual equilibrium monitoring probability μ^* to decrease and the equilibrium welfare level L^* to increase.*

This simple extension shows that the prospect of resource collapse can enhance incentives under certain conditions. We expect that the opposite can also be found under different assumptions about how actions within a period influence stock dynamics. Thus, analysis of a more general model could be fruitful. As noted already, some technical work must be done to ensure the existence of contractual equilibrium and to streamline the equilibrium characterization.

3 Interpretation and Implications

We next explore the implications of our model for structuring actual conservation agreements. Our model envisions a multi-period contract between an NGO and a local Community, established to achieve the NGO's goal of inducing the Community to preserve the resource while keeping the NGO's monitoring costs to a minimum. The contractual equilibrium is a grand contract that specifies the parties' beliefs and behavior across all periods and in every contingency. In an ideal world, perhaps the provisions would all be documented explicitly, but this is not necessary. In fact, the parties can coordinate on the equilibrium strategies without writing anything. Still, some level of documentation is typically needed in the real world, so that the parties can assure each other of their intentions and keep track of the details.

A particularly realistic approach to documentation is for the parties to express their intentions in a series of short-term contracts. In this form, the parties specify their behavior for the current period and agree on monetary rewards and punishments to take place in the next period, contingent on the outcome of the current period. They also agree, or at least expect, that the arrangement will be renewed in the next period if the parties behave as agreed currently. As an illustration, here is what such a short-term contract would specify:

- The Community agrees to protect the natural resource, and the NGO agrees to monitor at a specific level (μ, in our model, the probability of monitoring in any period).
- If the NGO does not monitor in the current period, then in the next period it must pay a prescribed amount to the Community, and the parties expect to renew the contract.[13]
- If the NGO monitors and there is no sign of exploitation, then the NGO must pay the specified amount, less the monitoring cost, and the parties expect to renew the contract in the next period. The Community implicitly shares in the cost of monitoring by receiving less when monitored than it does when not monitored.
- If the NGO monitors and obtains evidence that the Community exploited the resource, the parties may have specified that they should coordinate on protection and monitoring for sure until cooperation is restored. But then they renegotiate immediately, with the NGO making a much-reduced payment, so that they can renew the contract immediately and avoid the need for the NGO to monitor for sure.

As long as the equilibrium strategies are played (the Community protects the resource and the NGO monitors with the appropriate frequency), the parties agree to the same contract in each subsequent period. If the Community deviates and is caught (i.e., the Community exploits the resource and the NGO monitors and discovers this) or the NGO fails to make a specified payment, then the parties coordinate on the punishment actions but then renegotiate back to the desired short-term contract in the next period. To draw implications relevant for designing real-world conservation agreements, we elucidate these implications in terms of the short-term contracts.

[13] In the contractual equilibrium description, the NGO's payment is made during the renewal. If the NGO fails to pay or the parties otherwise fail to renew, they revert to exploitation and no monitoring until they renegotiate to restore cooperation.

Note that in the contractual equilibrium of our model, on the equilibrium path the NGO randomizes in its choice of whether to monitor in a period. The monitoring probability is just high enough to make the Community indifferent between protecting and exploiting the natural resource in a given period. Because the NGO is indifferent between monitoring and not (recall the monitoring cost is fully deducted from the payment to the Community), it is rational for the NGO to monitor with the correct frequency.[14] But since the Community does not observe the NGO's actual randomization device, one might wonder if the total joint payoffs could not be increased by a little less frequent monitoring. Indeed, with less monitoring, the Community's payoff would rise, given that in equilibrium it pays the cost of the monitoring. However, less frequent monitoring is not possible in equilibrium because the Community would then strictly prefer to exploit the resource.

We next provide a list of features, implications, and variations of the model that relate to the design and assessment of real conservation agreements. These are provided in three groups, having to do with the essential technological conditions for cooperation, necessary contractual ingredients, and aspects of optimal contractual arrangements.

Environmental Context and Essential Technological Conditions:

1. A resource exists that has value to the wider (global) community but is subject to the exploitation of a local Community. An NGO serves as a proxy for the global community.

2. The NGO has access to a monitoring technology sufficient to detect whether the Community exploits the resource. While the model assumes that the monitoring technology has a special structure, the practical requirement is that the technology is good enough to provide incentives to the Community.

3. The value of the resource to the NGO, net of the cost of monitoring the Community with great enough frequency, is greater than the Community's exploitation value.

4. The protection of the resource is a continuing activity, so if the resource is to be preserved in perpetuity then an arrangement between the Community and the NGO must be renewed regularly.

5. Both the NGO and the Community must have high enough discount factors (i.e. sufficiently low discount rates) for future benefits. In practice, this means that both

[14] If the Community is caught exploiting the resource then, during the punishment phase, the NGO monitors for sure and the Community must protect the resource; the payment is lower, but the Community has the incentive to protect because doing so allows cooperation to be restored.

the Community and the NGO must exhibit sufficient concern for the long term.

Features of a Successful Agreement to Preserve the Resource:

6. The agreement between the NGO and Community specifies conservation payments in exchange for protection of the resource. If the agreement is put in terms of a short-term contract, then the parties are expected to renew it conditional on satisfactory performance.

7. The Community must receive a sufficient share of the surplus that the agreement generates, so it must have sufficient bargaining power (some control over the renegotiation process). We expect to see some back-and-forth negotiations between the NGO and the Community in reality, rather than the NGO making a take-it-or-leave-it offer.

8. While an agreement to protect the resource can take the form of a short-term contract, the parties should anticipate how their agreements will be renegotiated over time.

9. The model assumes that monitoring is a binary choice, implying that randomization is optimal, but other monitoring technologies would perform similarly. Key features are that the level of monitoring can be observed by both parties and intermediate monitoring choices result in some uncertainty regarding detection.

Structure of an Optimal Agreement and the Equilibrium Strategies:

10. Since some monitoring of the Community is required to provide incentives to the Community, but monitoring is costly, the NGO and the Community have a shared interest in minimizing the amount of monitoring over time. Thus, under optimal monitoring, in a given period exploitation need not be detected with certainty.

11. An optimal contract specifies an up-front payment from the NGO and an additional payment in the following period if exploitation is not detected in the current period. If monitoring reveals that the Community has not protected the resource, then the agreement is not abandoned; however, the Community is "punished" by receiving a lower payment in the next period while it is still expected to protect the resource. If no further exploitation is uncovered by monitoring, then cooperation and full payments are restored.[15]

[15] Although our model has the NGO making a payment to the Community at the beginning of the period and rewards and punishments are levied in the subsequent period, an equivalent formulation would have the payment from the NGO to the Community occur at the end of each period. In fact, this framing of the timing of payments is the norm with real-world contracts, but does not signify a structural defect with such

12. Although the optimal contract specifies that the Community receive a lower payoff at the beginning of the period following one in which monitoring detected that it was not protecting the resource and that the NGO monitors for sure in the next period, this outcome is inefficient since monitoring costs are higher than necessary going forward. Because of this, both parties have an incentive in this case to renegotiate the contract resulting in a higher payoff for the Community resulting from reducing the probability of monitoring in the next period. [16]

13. In the model, since with an optimal contract the NGO must be indifferent between monitoring or not, the full cost of monitoring is deducted from the payment to the Community whenever monitoring occurs. In reality, the NGO may have a budget for the project and is willing to pay to the Community the full amount it has budgeted. However, the NGO cannot know if the Community is following the agreement without costly monitoring, so the cost of monitoring is deducted from the NGO's budgeted amount. Thus, it is in the Community's interest that monitoring be done as infrequently as possible while enabling the NGO to verify that the Community is indeed preserving the resource, and an optimal contract specifies precisely this level of monitoring.

14. In the real-world, seldom is the full monitoring cost deducted from the Community's payment conditional on the NGO's random monitoring choice. However, a straightforward variation of our model would have the NGO contracting with a third party to perform the monitoring at a particular level (i.e., with a specific frequency and with sufficient effort) for a fixed fee every period. If the third-party monitoring arrangement is verifiable, then this version of the model generates the same results as we have derived herein. With third-party monitoring, the payment to the Community would not vary with monitoring as long as exploitation is not detected.

These implications of the model suggest how to analyze individual cases, which we do in the next section.

contract and is not in conflict with our model. The important feature of an optimal contract is that *if* the Community fails to protect the resource, and this is discovered by NGO monitoring, then a "punishment" (i.e. lower payment) must occur, whether it is paid at the end of the current period or at the beginning of the next period.

[16] The payoffs specified in the optimal contract, for the case when the Community is detected exploiting the resource, are required to provide the correct incentives in the initial period and define the default values for the renegotiation bargaining.

4 Case Studies

In this section we present three actual cases of conservation contracts. We use our model to evaluate their varying degrees of success and to suggest how future conservation agreements may be structured to avoid some of the problems that the cases illustrate.

4.1 Forest protection in Cambodia[17]

A conservation agreement was developed in Chumnoab Commune, Cambodia to maintain and protect the forest, wildlife, and crocodile habitat and to assist in efforts to combat illegal hunting and wildlife trade within the Commune. The parties to the agreement were the Commune Council of Chumnoab Commune (the "Council") and the NGO Conservation International Cambodia ("CI"). The agreement was endorsed by the District Governor and the District Police chief. The initial agreement covered the year from May 25th, 2006, to May 24th, 2007, with the understanding that the terms would be renewed on a yearly basis indefinitely. It was agreed that the parties could transition to a long-term agreement, which would involve reviewing terms periodically but not necessarily every year.

CI and the community engaged in a participatory land use planning process to map the areas of the Commune and decide which activities may or may not occur in which areas. This is recorded in a Participatory Land Use Plan (PLUP). The commune members agreed not to engage in the setting of snares to capture wildlife, not to bring dogs into forest areas, to follow crocodile protection rules set forth in the agreement, and to inform the Chumnoab Commune Natural Resource Management Council (CNRMC) of any observed and rumored hunting and wildlife trade activities. One of the main threats is from clearing forest for rice production. Community rangers agreed to patrol forest areas, remove snares, and report observed hunting and wildlife trade activities to the CNRMC. The CNRMC agreed to inform CI and the Forestry Administration of any violations, including clearing.

In return, CI agreed to provide eight water buffalo to the villagers of the Chumnoab commune ($4,000 USD total purchase cost) at the end of the first agreement period (one year). In addition, CI agreed to (a) transfer $25 USD per month to the resident teacher at the Chumnoab public school; (b) provide financial support for the construction of a school building ($500 USD); (c) provide funding for patrolling activities of the community rangers and police support (approximately $3,600 USD for 15 days of patrol per month by four people); and (d) transfer $30 USD per month to the CNRMC for organizing Community Ranger patrols, collecting and submitting patrol reports, informing all Commune members

[17] The following references provide some background information for the case study discussed in this section: Conservation International (2007); Milne and Niesten (2009); Ouk

of the terms of the agreement, and facilitating effective, transparent and equitable delivery of benefits. Summing up, approximately $8,760 USD was to be spent annually to protect 6555.42 hectares of forest.

Compliance monitoring is conducted primarily by the Forestry Administration, an independent party. The Forestry Administration conducts patrols jointly with community rangers, using a transect monitoring regime. One of the main observable variables is whether forest was cleared. Non-compliance with the agreement is to be reported immediately to Conservation International. Community rangers from the commune, in coordination with the local police, are supposed to ensure that land and resource use within the commune area complies with the terms of the agreement. Those found in violation of the agreement are to be reported to the relevant authorities or to the community committee, depending on the violation. The sanctions specified for violating the terms of the agreement are shown in Table 1. In cases where a family violates the agreement and loses a water buffalo, this animal is to be given to another family on the list.

Table 1. Transgressions and sanctions in the Chumnoab Agreement

Transgressions	Sanctions
1–2 families with water buffalo violate PLUP rules	Families lose water buffalo, and commune receives warning of 50% reduction of benefit package in the following year.
3 or more families with water buffalo violate PLUP rules	Families lose water buffalo, and commune benefit package for the subsequent year reduced by 50%
1–2 families without buffalo violate the PLUP rules	These families go to bottom of list for receiving water buffalo, and commune receives warning of 50% reduction of benefit package in the subsequent year.
3 or more families without water buffalo violate the PLUP rules	These families go to bottom of list for receiving water buffalo, and the commune benefit package for the subsequent year is reduced by 50%.

During the initial agreement period, a violation of the contract occurred, whereby community members cleared approximately 12 hectares (ha) of forest. The community initially claimed that the infraction was not its fault, as the boundaries had not been clearly marked, but the community ultimately conceded that the agreement had been violated (warranting a sanction). Rather than go through with the sanction prescribed by the agreement, the parties renegotiated in a way that benefited both the NGO and the community relative to what would have happened under the sanction. They agreed to a one-time waiver of the penalty clause, and to proceed with the following revised stipulations instead:

1. Before the current agreement can enter into effect, the Council will provide to CI a list of names of the people responsible for the clearing, which will also indicate whether they received buffalos, and whether they participated in community ranger patrols during the period of the first agreement;

2. CI and the CNRMC will designate an additional, previously unprotected, 12 ha elsewhere for protection, to substitute for the 12 ha that were cleared in violation of the agreement (the new areas protected are suitable according to the land use plan);

3. The families responsible for the clearing will be allowed to cultivate the cleared land for one season, after which the area will revert to protected status;

4. The community benefit package will be reduced for one year from 4 buffalo to 2;

5. With regard to the current agreement, CI will not be disposed toward similar flexibility with respect to sanctions in the event that further violations take place.

A second agreement was entered into from May 25th 2007 to May 24th 2008. When the parties complied with this agreement, CI renegotiated a new agreement with the commune for the following year (May 25th, 2008, to May 24th, 2009). The parties complied with the agreement in the third year as well. In 2009, land reform affected the area, which has led to changes in the project and rendered the former agreement moot.

This case illustrates many of the key features of our model. Regarding the environmental context and essential technical conditions, note that CI was willing to make payments that appear adequate to induce the Community to protect, so we know the resource has value to the NGO that exceeds the Community's exploitation value. The protection of the resource is a continuing activity and the agreement between CI and the Community explicitly envisions yearly renewal. The contract specifies payments in exchange for conservation effort and it describes how the Community will be punished in the event of noncompliance, consistent with our model. Further, the structure of the agreement indicates that both the NGO and the community have reasonably high discount factors.

The model assumes that both parties exercise bargaining power, both initially and in every renegotiation. In the Cambodia case, the parties settled the land clearing issue through an ex-post investigation and discussion, in which the Community agreed to put additional forest area under protection in exchange for a reduction in punishment severity.

The resolution thus involved both parties giving up something, an indication that both have bargaining power.

Note that in this case study an independent party monitors compliance with the participation of community members every month or so. There are more intensive, daily patrols around dragonfish ponds during the breeding season. Once a year there is a full assessment by a larger research team. The level of monitoring effort in practice fluctuated, particularly for the full annual assessment. A key feature of our model is that since monitoring is costly, an optimal contract will not involve constant monitoring, but will reduce the amount of monitoring to some extent. But without knowing exactly by how much CI values the conservation activity and also knowing the Community's exploitation value, it is not possible to determine whether the monitoring is excessive under this agreement. Our theory would suggest however, that, if the contract were to continue, savings from a reduction in the amount of monitoring might be something over which to negotiate.

Although our model specifies that the full cost of monitoring is deducted from the payment to the community whenever monitoring occurs, this is not an explicit feature of this contract. However, as discussed above in Section 3, there is an alternative interpretation of our model involving third party, independent, monitoring that is consistent with optimal conservation agreements, but with constant average monitoring costs being paid to a third party monitor on a period-by-period long-term contract. In fact, the monitoring specified in this Cambodian forest protection agreement involves precisely such a third party monitor, the Forestry Administration. Alternatively, since monitoring is evidently done every period in this case, one could just as well view the payment to the commune as a payment net of the costs of monitoring, which are paid directly to the Forestry Administration for monitoring.[18]

It is interesting to note that in the first period of this agreement, there were, in fact, violations. However, instead of administering the penalties initially specified, which apparently would have reduced joint welfare, the parties renegotiated to less severe sanctions, as expected. While the model does not predict violations on the equilibrium path, it predicts that the contract will be renegotiated if there were detected violations, as occurred in the first period here.

Taking all things under consideration, we conclude that this particular agreement corresponds well with our model and theory of an optimal conservation agreement. The structure of the contract accords with our model of an optimal contract, the outcomes are consistent with our theory, and, most important, the contract was largely successful in achieving its

[18]It also should be noted that this agreement includes some monitoring by community patrols that are paid for by CI. Although self-monitoring by the community is not addressed by our model, if effective it can help further reduce the cost of monitoring and thereby benefit both parties.

intended goals of protecting the forest.

4.2 Laos deer conservation[19]

The Lower Mekong Dry Forests in Southeast Asia have been identified as a global priority for biodiversity conservation. The dry forest area in Savanakhet Province in Lao PDR is some of the last remaining habitat for the endangered Eld's deer (*Cervus eldii*). In 2003, the Wildlife Conservation Society (WCS) and Smithsonian Institution initiated a conservation-payments program in Laos with the aim of reducing threats and increasing the size of the Eld's deer population. Hunting and habitat clearing by villagers living nearby was threatening the deer, despite legal protection by a wildlife conservation law since 1995.

During an initial workshop, the NGOs and villagers (the "Community") discussed the threats to the deer and how to address poaching and habitat destruction. The Community agreed to establish community patrols that would report and stop poachers. The Community also promised to maintain habitat and conserve resources in other ways, such as by not expanding rice paddies and by keeping cattle out of water holes. In exchange, the NGOs agreed to pay an annual cash incentive to each of three villages located near the deer population. The agreement stated that the NGOs would return at the end of each year to assess the deer population and would make the payments only if the deer population had increased. Villagers decided to use the payments for a village development fund and to pay per diems (approximately 2 USD/person/day of activity) for meetings, patrolling, and education work by the Village Conservation Team (VCT). The payment was initially $300 USD, and was increased to $450 USD in the second and third years to be able to fund complete village development projects.

WCS, SI, government staff, and villagers monitored the deer population size by transect lines and surveyed habitat to see whether area used by deer had increased or decreased. In addition, villagers were asked to report all deer sightings to a literate person in their village to record and submit the data to WCS. According to WCS, shortage of funds and WCS staff resources prohibited a more rigorous monitoring methodology. Because there were so few deer, it was difficult to actually estimate the population. So the NGOs would assess whether the deer seemed to be using more habitat (i.e. either there were more deer or they were accessing more habitat, either of which is a good sign).

At the end of the first year, monitoring indicated that there was no change in the deer population, but even so the NGO decided to make the payment to the Community because

[19] The following references provide some background information for the case study discussed in this section: McShea (2015); Svadlenak-Gomez, Clements, Foley, Kazakov, Lewis, Miguelle, and R.Stenhouse (2007)

villagers expressed excitement and support for the program. At the end of the second year, monitoring indicated a decrease in the deer population and there was encroachment by villagers' rice paddies in the deer habitat. Nonetheless, the NGO made the payment to the Community because of fear of a lack of future cooperation, and because the NGO was reluctant to deny the Community a payment for much-needed school expenses.

Shortly thereafter, the NGO decided to abandon the agreement, citing doubts about the merits of the program and limited funding. The deer project is now run by a different NGO (WWF) and is focused on land-use planning, villager-led patrolling of the sanctuary, and the development of sustainable livelihood opportunities.

The agreement of this case is one that appears to follow the basic design of an optimal contract of our model, but failed due to some key elements of an optimal contract and also from an evident unwillingness or inability to implement the punishment terms after the Community violated the protection terms of the agreement. In addition, the NGO did not try to renegotiate an alternative punishment with the Community, as we saw above for the similar case of forest protection in Cambodia.

As with the Cambodian case, all the major conditions of our model were satisfied—there was sufficient value from preservation of the Eld's deer for NGOs to assemble funds to compensate the Community for foregoing poaching and protecting habitat. Further, it was recognized by both sides that this was a project that would have to continue into the indefinite future for it to be deemed a success; that is, both the NGOs and the Community clearly went into the initial agreement with an understanding that their arrangement would continue into future periods. So, what can we identify as reasons why this agreement failed?

From the history of how the agreement failed we can identify several key flaws. Probably the most important one was that insufficient attention was devoted to specifying an adequate monitoring regime. The NGO relied on a vague plan to have a yearly assessment of the deer population and judge the compliance of the Community by whether or not this assessment showed an increase in the deer population. As they discovered, it was not a simple task to assess the deer population with sufficient accuracy to justify enacting the "punishment" of the agreement if the population declined. Perhaps better (i.e. more costly) methods for assessing the deer population could have been used, but in any case, *ex post* it appears insufficient attention was given to this key monitoring requirement of the agreement.

Another major flaw with the agreement was a failure of the NGOs to anticipate that they would be unwilling to actually implement the punishment (i.e. a withholding of the full payment) if the Community failed to comply with the agreement and they were unable to renegotiate the agreement, as would be expected according to our model. Even when there was a clear violation of the agreement, as, for example, in its second year when rice paddy encroachment into the forest deer habitat was observed by all, the NGO was unwilling to withhold the payment as required by a strict enforcement of the agreement's provisions. A possible reason for the NGOs reluctance is that the agreement itself did not appropriately specify what should happen when the Community is caught exploiting and the parties fail to renegotiate. The model suggests that in such an even the NGO should monitor much more closely while the Community protects. Anticipating that such close monitoring is wasteful, the parties should then renegotiate to normal levels of monitoring but with the NGO paying only a small amount to the Community. It seems the NGO was not willing to monitor closely even after a violation and a disagreement.

We suggest that an examination of the reasons for the agreement's failure validates the conditions and requirement of the model. Both the inadequate monitoring specified by the agreement, the unwillingness of the NGO to fulfill the terms of the contract, and an inability to establish renegotiation were failures of necessary conditions for our model to apply.

4.3 Grey whale habitat protection in Mexico[20]

Laguna San Ignacio is situated on the Pacific Coast of Baja California Sur, Mexico. It is the world's last untouched breeding ground for Pacific gray whales; more than half of the world's gray whales calves are born inside Laguna San Ignacio and the neighboring lagoon of Ojo de Liebre. In addition to providing grey whale habitat, Laguna San Ignacio hosts at least 221 other animal species, including numerous birds, green sea turtles, and bottlenose dolphins.

In 1994, Mitsubishi proposed to establish a salt plant at Laguna San Ignacio, but due to local and international pressure, the project was shelved in 2000. To conserve the area over the long term and prevent future coastal development, Mexican NGO Pronatura suggested the option of an easement. A conservation easement is a voluntary, legally binding agreement between two parties in which the land use rights of one party are restricted, with the objective of preserving in perpetuity natural resources, scenic beauty, or historical and cultural values of the land. In 2005, the Laguna San Ignacio Conservation Alliance established a 120,000-acre conservation easement comprising all the communal lands within the Ejido Luis Echeverria Alvarez on the southern shore of Laguna San Ignacio.

There are four parties to the agreement, each with a specific role. The Ejido Luis Echeverria (hereafter called "the Community") agrees to limit coastal development. Pronatura (the "NGO") monitors compliance. The International Community Foundation (ICF) is a San Diego foundation responsible for disbursing funds to the Community. ICF maintains

[20]The following references provide some background information for the case study discussed in this section: Alarcon; Gjertsen and Niesten (2010); Lopez; McEnany; Youatt; **?**

a trust fund and manages it as a third party so there is transparency and accountability. Maijanu is an organization that was created in the Community to receive and manage the funds disbursed through the easement.

Pronatura conducts bi-annual monitoring of the area to determine compliance with the terms of the easement. A team of biologists, GIS experts, and lawyers from Pronatura visit the same sites every six months and take photos and compare to originals. They also interview 10-15 community members about whether they have noticed any changes. Community members also monitor throughout the year. ICF maintains a fund that disburses approximately $15,000 USD per year to Pronatura for monitoring.

Each year, if Pronatura determines that the Community has met its obligations, ICF is supposed to disburse to Maijanu the annual interest generated from the Ejido Luis Echeverria Alvarez Seed Fund, which was capitalized in the amount of $650,000 USD. These annual payments amount to approximately $25,000 USD per year. ICF had planned for an increase in payments over time, but the Community has chosen to maintain a flat annual $25,000. As a result, the fund has grown to $808,000 USD. The Community chose to use the payments for community projects rather than divide the funds as individual payments to members. The payments can be used for any community development projects that are not harmful to the environment and that do not contradict the terms of the contract. Every year any member can present a project proposal that will be reviewed by the community leadership and all the members vote in a general assembly for the proposals.

According to the agreement, if the Community's obligations in the contract are not met, then the community payments will not be disbursed. If the violation created damage that can be restored, then the payments may be restarted once the damage is restored. If the damage can not be restored, the payments will be halted permanently. Since the contract is signed in perpetuity, compliance is required each and every year. When compliance is lacking, not only can the payments be halted, but Pronatura can also take legal action to force compliance, which could include cessation of the illegal activity and restoration.

Thus far, the terms of the easement have been met every year by the Community, and they have received the community payments every year.

The Ejido Luis Echeverria conservation easement is a contract that closely matches the design of our model. The grey whale habitat clearly has a high value to the public, given that NGOs were able to raise millions of dollars for its protection (an indication of the willingness to pay). This allows the NGO to pay the Community to forego future destructive development projects. The funds more than cover the annual payments to the Community and the monitoring costs. Preserving the habitat requires ongoing effort, which the parties clearly recognized by forming a contract in perpetuity. The NGO and the Community both appear to have high discount factors. The Community is accepting very low annual payments, compared to what it might be able to earn by selling its land. The NGO tied up a great deal of money in a trust fund that is to be used for long-term conservation.

On bargaining power, note that the Community has been accepting a fairly low monetary amount over time, and has not attempted to increase the annual payments or renegotiate contract terms. Rather than this being due to low bargaining power, we think it is because the Community interests are mostly aligned with the NGO; that is, the Community receives value from choosing to protect the habitat, due to tourism and fishing opportunities and the interest in maintaining a simple lifestyle. However, this may change with future land speculation, particularly with a paved road and electricity due to reach the Community imminently. Thus, as the fundamentals change, the agreement will encounter stress and we predict that renegotiation will occur.

The key is that the payment, as well as behavior in subsequent interaction, is conditioned on the outcome of monitoring. The contract specifies that if monitoring reveals that the Community has protected the resource, then the Community will receive the same payment in next period, and so on into the future. If monitoring reveals that the Community has not protected the resource, then the Community will receive the payments only after reversing the damage from exploitation. The contract also states that payments will be halted if the damage cannot be restored, which we interpret as disengagement (some degree of exploitation and no payments) unless and until the parties choose to renegotiate. Our analysis anticipates that if unrestorable damage did occur, the parties should nonetheless find it optimal to renegotiate, in such a way that would punish the community while rewarding the NGO.[21]

Consistent with the model, the agreement specifies the monitoring activity in detail and accounts for its cost. In fact, the agreement is quite monitoring-intensive. Monitoring occurs at specified intervals and does not vary a great deal. However, it does involve some minimization of costs, as monitoring could occur more frequently or could involve more detailed site visits (interviewing more community members, inspecting all land, etc). The model specifies that the full cost of monitoring is deducted from the payment to the Community whenever monitoring occurs, which can be interpreted as the case here, as monitoring costs are deducted from a separate account.

We regard the Laguna San Ignacio agreement, along with the other case studies, as confirming the message of our modeling exercise regard the ingredients essential for cooperation. However, the Laguna San Ignacio agreement has a potentially important element that is

[21] For instance, they could agree to allow Pronatura to deduct a penalty amount from the trust fund, to spend on conservation efforts elsewhere.

outside our repeated-game model: some degree of external enforcement that may enhance incentives to cooperate, beyond what could be achieved by self-enforcement alone. Some aspects of the contract may be enforceable in Mexican courts. In particular, Pronatura (with perhaps enhanced standing as a Mexican organization) can take legal action to force the Community to cease illegal activity and to restore damage. The interaction between self-enforcement and external enforcement is a topic ripe for further study, both theoretically and empirically, and is on our research agenda going forward.[22]

5 Conclusion

Our results suggest several implications for the design and implementation of conservation agreements:

1. Reaching an agreement requires that it generate enough surplus for the parties to share.

2. The Community must obtain a sufficient share of the surplus from each period, so it is important for the Community to have adequate bargaining power (i.e., some control over the renegotiation process).

3. The NGO and Community must have high enough concern for the future to achieve protection in equilibrium.

4. Parties should anticipate how their agreements will be renegotiated over time, in particular following any infraction. For example, opportunities for renegotiated punishments and monitoring probabilities following infractions can increase total joint payoffs by reducing the costs of monitoring and also provide greater incentives for cooperation.

5. Since monitoring is inherently costly and thus reduces the total payoff available for division between the two parties, they both have a shared interest in minimizing the amount of monitoring over time to a level just sufficient to provide the necessary incentives to preserve the resource.

We think it would be worthwhile to continue the theoretical exploration by looking at more general productive and monitoring technologies, outside options (in particular for the NGO), and resources with growth-depletion dynamics. Furthermore, as in the case study

[22]Theorists, including some of the coauthors of this paper, are currently developing models of the interaction between external and self-enforcement. A recent example is **?**.

from Mexico, the combination and interaction of self-enforcement and external enforcement is an important topic for continued research. On the applied side, it would be useful to look carefully at specific settings beyond conservation agreements, such as REDD (Reducing Emissions from Deforestation and Forest Degradation) contracts. Finally, we think that our treatment of the Community as single player is likely hiding many interesting issues on the relation between the incentives of individuals within the Community and the Community as a whole, including with regard to how bargaining takes place.

A Appendix

This appendix proves the theorems in the paper, by applying the Miller-Watson algorithm (Miller and Watson 2013) that characterizes the unique contractual equilibrium value (CEV) set V.[23]

We begin by proving that the stage game action profile taken under disagreement should be ER when the NGO is being punished, and PM when the Community is being punished. Since the CEV set takes the form $V = \{z^C + (\zeta, -\zeta) \mid \zeta \in [0, d]\}$, the players' expected payoffs starting in the action phase, up to a lump sum, can be expressed as $w(\alpha, \eta) \equiv (w_C(\alpha, \eta), w_N(\alpha, \eta))$, where $\alpha \in \Delta A$ is a mixed stage game action profile, and $\eta : S \to \mathbb{R}$ describes the value taken from the NGO and given to the Community (constraints on η are addressed below), starting from the next period, as a function of the realized signal:

$$w_C(\alpha, \eta) \equiv u_C(\alpha) + \delta_C \mathbb{E}\left[\eta(s)|\alpha\right],$$
$$w_N(\alpha, \eta) \equiv u_N(\alpha) - \delta_N \mathbb{E}\left[\eta(s)|\alpha\right].$$

The following game matrix defines $w(a, \eta)$ for this game:

	M	R
P	$\delta_C\eta(MG),\ (b-c) - \delta_N\eta(MG)$	$\delta_C\eta(RG),\ b - \delta_N\eta(RG)$
E	$e + \delta_C\big(\lambda\eta(MB) + (1-\lambda)\eta(MG)\big),$ $-c - \delta_N\big(\lambda\eta(MB) + (1-\lambda)\eta(MG)\big)$	$e + \delta_C\eta(RG),\ -\delta_N\eta(RG)$

[23] Note that Miller and Watson (2013) address the case of heterogeneous discount factors in their Appendix, where variables expressed in total discounted utility terms are marked with tildes (e.g., $\tilde{V}$) to distinguish them from variables expressed in average utility terms. Here, since all variables are expressed in total utility terms throughout, we do not use tildes.

State N, disagreement To find the optimal disagreement action profile for state N, the Miller-Watson algorithm solves the following problem, for any payoff span $d \geq 0$:

$$\gamma^N(d) \equiv \max_{\eta,\alpha} \frac{\pi_N u_C(\alpha) - \pi_C u_N(\alpha) + \psi \mathbb{E}\left[\eta(s)|\alpha\right]}{1-\psi}, \quad \text{s.t.} \begin{cases} \eta : S \to \left[-d, 0\right], \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \langle A, w(\cdot,\eta)\rangle, \end{cases} \tag{14}$$

where ψ is the "weighted" discount factor $\psi \equiv \pi_C \delta_N + \pi_N \delta_C$.

Lemma 1. *The stage game action profile $\alpha^N = ER$, combined with zero transfers ($\eta^N(s) = 0$ for all s), solves (14). The maximized value is $\gamma^N(d) = \frac{\pi_N}{1-\psi} e$.*

Proof. Because ER is a Nash equilibrium of the stage game, α^N and η^N satisfy the constraints of the maximization problem. With these selections, the objective function attains the value $\frac{\pi_N}{1-\psi} e$.

Next consider any action profile in which the Community selects E and the NGO plays M with probability μ and R with probability $1-\mu$, where $\mu > 0$. When the NGO selects M, the public signal is MB with probability λ and MG with probability $1-\lambda$. In order for the NGO to have the incentive to select M with positive probability, we need:

$$(-c) + \delta_N \left(\lambda(-\eta(MB)) + (1-\lambda)(-\eta(MG))\right) \geq 0 + \delta_N(-\eta(RG)). \tag{15}$$

The value of the objective function is thus $\frac{1}{1-\psi}\big(\pi_N e - \mu\pi_C(-c) + \psi\mathbb{E}\left[\eta(s)|E,\mu\right]\big)$. Note that $\mathbb{E}\left[\eta(s)|E,\mu\right] = \mu\lambda\eta(MB) + \mu(1-\lambda)\eta(MG) + (1-\mu)\eta(RG)$. Since $\eta(s) \in \left[-d, 0\right]$ is required, Eq. 15 implies that $\lambda\eta(MB) + (1-\lambda)\eta(MG) \leq -\frac{c}{\delta_N}$, which further implies that $\mu\lambda\eta(MB) + \mu(1-\lambda)\eta(MG) + (1-\mu)\eta(RG) \leq -\frac{\mu c}{\delta_N}$. Thus the value of the objective function does not exceed $\frac{1}{1-\psi}\big(\pi_N e + \mu c\big(\pi_C - \frac{\psi}{\delta_N}\big)\big) \leq \frac{\pi_N}{1-\psi} e$.

It is easy to see that other action profiles—those in which the Community selects P with positive probability—lead to even lower values of the objective function. Thus, we conclude that α^N and η^N solve the maximization problem that defines γ^N, and the resulting value of γ^N is as stated. □

State C, disagreement To find the disagreement action profile for state C, the Miller-Watson algorithm solves the following problem, for any payoff span $d \geq 0$:

$$\gamma^C(d) \equiv \min_{\eta,\alpha} \frac{\pi_N u_C(\alpha) - \pi_C u_N(\alpha) + \psi \mathbb{E}\left[\eta(s)|\alpha\right]}{1-\psi}, \quad \text{s.t.} \begin{cases} \eta : S \to \left[0, d\right], \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \langle A, w(\cdot,\eta)\rangle. \end{cases} \tag{16}$$

Lemma 2. *Suppose that Eq. 13 holds. Then the stage game action profile $\alpha^C = PM$, combined* *with* $\eta^C(MB) = 0$, $\eta^C(MG) = \frac{e}{\delta_C \lambda}$, *and* $\eta^C(RG) = \frac{c}{\delta_N} + \frac{e}{\delta_C \lambda}$, *solves (16). The minimized value is*

$$\gamma^C(d) = \frac{-\pi_C(b-c) + \frac{\psi e}{\delta_C \lambda}}{1-\psi}.$$

On the other hand, if (13) does not hold then $\gamma^C(d) = \frac{\pi_N}{1-\psi} e = \gamma^N(d)$.

Proof. First note that PR cannot be supported because the Community has no incentive to choose P when the NGO does not monitor. Also note that play of E yields a high value of the objective function relative to specifying that P is to be chosen. Thus, let us look at stage-game action profiles in which the Community selects P and the NGO chooses M with probability μ and R with probability $1-\mu$, where $\mu > 0$.

In order for the NGO to have the incentive to select M with positive probability, we require $(b-c) + \delta_N(-\eta(MG)) \geq b + \delta_N(-\eta(RG))$, which simplifies to

$$\eta(RG) \geq \eta(MG) + \frac{c}{\delta_N}. \tag{17}$$

Likewise, in order for the Community to have the incentive to choose P, we require $0 + \delta_C[\mu\eta(MG) + (1-\mu)\eta(RG)] \geq e + \delta_C[\mu\lambda\eta(MB) + \mu(1-\lambda)\eta(MG) + (1-\mu)\eta(RG)]$, which simplifies to

$$\eta(MG) \geq \eta(MB) + \frac{e}{\delta_C \mu \lambda}. \tag{18}$$

With the specified stage-game action profile, the value of the objective function is $\frac{1}{1-\psi}\big(-\pi_C(b-\mu c) + \psi\mathbb{E}[\eta(s)|P,\mu]\big)$. Note that $\mathbb{E}[\eta(s)|P,\mu] = \mu\eta(MG) + (1-\mu)\eta(RG)$. Since this value is increasing in $\eta(MG)$ and $\eta(RG)$, and since we have a minimization problem with $\eta(s) \in \big[0, d\big]$, it is optimal to have Eq. 17 and 18 hold with equality and to have $\eta(MB) = 0$. Thus $\eta(MB) = 0$, $\eta(MG) = \frac{e}{\delta_C \mu \lambda}$, and $\eta(RG) = \frac{c}{\delta_N} + \frac{e}{\delta_C \mu \lambda}$. Plugging these values into the objective function, we obtain the value

$$\frac{-\pi_C(b-\mu c) + \psi\left(\frac{e}{\delta_C \mu \lambda} + (1-\mu)\frac{c}{\delta_N}\right)}{1-\psi}.$$

This value is clearly decreasing in μ, so it is optimal to have $\mu = 1$ and we get the expression for γ^C shown in the statement of the lemma. Note that the condition on d is required for the chosen values of $\eta(MB)$, $\eta(MG)$, and $\eta(RG)$ to be feasible. □

Agreement

Proof of Theorem 1 *on p. 14*. The next step in the Miller-Watson algorithm is to calculate the maximal fixed point of the function $\Gamma \equiv \gamma^N - \gamma^C$. From Lemmas 1 and 2, we see that

$$\Gamma(d) = \frac{1}{1-\psi} \cdot \begin{cases} \pi_N e + \pi_C(b-c) - \psi\frac{e}{\delta_C \lambda} & \text{if (13) holds} \\ 0 & \text{otherwise.} \end{cases} \tag{19}$$

The theorem considers only the case in which Eq. 13 holds; the first line of Eq. 19 at its highest fixed point is guaranteed to be greater than the second in this case. (Sufficiently high $\psi < 1$ is sufficient for Eq. 13 to be satisfied, so the case is not empty.) Assuming Eq. 13, we can then write the maximal fixed point of Γ as:

$$d^* = \frac{\pi_N e + \pi_C(b-c) - \psi \frac{e}{\delta_C \lambda}}{1-\psi} \tag{20}$$

This number d^* is the payoff span of the CEV set.

We next determine the welfare level of the CEV set, L^*, which is the greatest joint value that can be supported when the span of continuation payoffs from the next period is d^*. The Miller-Watson algorithm does this by first calculating:

$$\begin{aligned} \rho(d^*) \equiv \max_{\eta,\alpha} \; & u_C(\alpha) + u_N(\alpha) + (\delta_C - \delta_N)\mathbb{E}_s\left[\eta(s)|\alpha\right], \\ \text{s.t.} \; & \begin{cases} \eta : S \to [-d^*, 0], \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \langle A, w(\cdot,\eta)\rangle. \end{cases} \end{aligned} \tag{21}$$

Since we assume Eq. 13 holds, we know from Lemma 2 that PM is enforceable on span d^*; the question is whether it can be improved upon by reducing the probability of monitoring.[24] Since the stage-game optimum PR cannot be enforced by any η, we conclude that, in order for the Community to play P, the NGO must play M with some probability $\mu > 0$. Letting (P, μ) indicate this action profile, Eq. 21 becomes:

$$\begin{aligned} \rho(d^*) \equiv \max_{\eta,\mu} \; & b - \mu c + (\delta_C - \delta_N)\big(\mu\eta(MG) + (1-\mu)\eta(RG)\big), \\ \text{s.t.} \; & \begin{cases} \eta : S \to [-d^*, 0], \\ (P,\mu) \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \langle A, w(\cdot,\eta)\rangle. \end{cases} \end{aligned} \tag{22}$$

Recalling the analysis from the proof of Lemma 2, we see that Eq. 17 is required for the NGO to choose M with positive probability. Furthermore, Eq. 18 is required for the Community to select P. For any given $\mu > 0$, it is optimal for these constraints to bind regardless of the sign of $\delta_C - \delta_N$. Consider two cases.

1. $\delta_C \geq \delta_N$: Then for any μ it is optimal to maximize $\eta(RG)$ and $\eta(MG)$ subject to Eq. 17 and Eq. 18 and $\eta : S \to [-d^*, 0]$. This is solved at $\eta(RG) = 0$ and $\eta(MG) = -c/\delta_N$. It follows that μ should be minimized subject to $\eta(MB) \geq -d^*$. The lowest value of μ consistent with

[24]When $|\delta_N - \delta_C|$ is very large the solution to Eq. 21 may be to play ER and choose η fixed at either 0 or $-d^*$; in this case each player receives its stage game minimax payoff along the equilibrium path, which implies that the span is zero and there must never be any transfers. Since PM is enforceable under Eq. 13, the players can certainly do better.

these conditions is that μ^* which solves $\frac{c}{\delta_N} + \frac{e}{\delta_C \mu^* \lambda} = d^*$, which simplifies to

$$\mu^* = \frac{\delta_N e}{\lambda \delta_C \big(\delta_N d^* - c\big)}. \tag{23}$$

It is straightforward to show that $\mu^* = 1$ when Eq. 13 is satisfied with equality, and that μ^* is decreasing in d^* but always strictly positive.

2. $\delta_C < \delta_N$: Then for any μ it is optimal to minimize $\eta(RG)$ and $\eta(MG)$ subject to Eq. 17 and Eq. 18 and $\eta : S \to [-d^*, 0]$. This is solved at $\eta(RG) = -d^* + \frac{e}{\delta_C \mu \lambda} + \frac{c}{\delta_N}$ and $\eta(MG) = -d^* + \frac{e}{\delta_C \mu \lambda}$. In this case Eq. 22 simplifies to

$$\max_{\mu}\ b - \mu c + (\delta_C - \delta_N)\Big(-d^* + \frac{e}{\delta_C \mu \lambda} + (1-\mu)\frac{c}{\delta_N}\Big). \tag{24}$$

Because $\delta_C < \delta_N$ and the term in large parentheses is decreasing in μ, this problem may have an interior solution or a corner solution, depending on the parameters. If $\delta_N - \delta_C > 0$ is extremely large, the corner solution $\mu^* = 1$ that maximizes μ may arise. If $\delta_N - \delta_C > 0$ is small, the corner solution above in Eq. 23 that minimizes μ may arise.

The Miller-Watson algorithm identifies L^* and the endpoints of the CEV set in terms of d^*, using the functions γ^C and ρ. Note that the level is not simply the discounted value playing (P, μ) in every period along the equilibrium path; since the players have different discount factors, the welfare level must also account for the utility gained or lost through the transfers. The endpoint favoring the NGO is given by:

$$z_C^{C*} = \frac{1}{1-\delta_C}\big(\pi_C \rho(d^*) + (1-\psi)\gamma^C(d^*)\big)$$
$$z_N^{C*} = \frac{1}{1-\delta_N}\big(\pi_N \rho(d^*) - (1-\psi)\gamma^C(d^*)\big).$$

We then have $z^{N*} = z^{C*} + (d^*, -d^*)$ and $L^* = z_C^{C*} + z_N^{C*} = z_C^{N*} + z_N^{N*}$. □

Comparative statics

Proof of Theorem 2 on p. 16. Taking partial derivatives and imposing our parametric assumptions yields the comparative statics in the theorem. Detailed computations are available on request. □

The case of equal discount factors

Proof of Theorem 3 on p. 16. In the special case of equal discount factors, i.e., $\delta_C = \delta_N \equiv \delta$, the payoff span of the CEV set is given in average terms as:

$$d^* = \begin{cases} \pi_C(b-c) + \pi_N e - \frac{e}{\lambda} & \text{if (13) holds,} \\ 0 & \text{otherwise.} \end{cases}$$

Assume that Eq. 13 holds. Then the maximization problem defining ρ simplifies to maximizing $b - \mu c$ over action profiles that are enforced using an average payoff span of d^*. The solution is to have the Community select P and have the NGO choose M with probability $\mu^* = \frac{e}{\lambda(\frac{\delta}{1-\delta}d^* - c)}$, so we have $\rho(d^*) = b - \mu^* c$, $L^* = b - \mu^* c$, and $\gamma^C = \pi_N e + \pi_C(b - c) - \frac{e}{\lambda}$. The welfare level L^* is just the discounted value of receiving $b - \mu^* c$ in every period. The comparative statics follow from taking derivatives. □

The simple stock extension

Proof of Theorem 4 on p. 20. For the simple stock extension, some algebraic manipulation reveals that

$$z_C^N = \pi_C L + \frac{\pi_N e}{1 - \delta\beta}, \; z_C^C = \frac{e}{(1-\delta)\lambda} + \pi_C L - \frac{\pi_C(b-c)}{1-\delta}, \text{ and}$$

$$d^* = z_C^N - z_C^C = \frac{\pi_N e}{1-\delta\beta} + \frac{\pi_C(b-c)}{1-\delta} - \frac{e}{(1-\delta)\lambda}.$$

The effect on μ^* of a small increase in β is given by $\frac{\partial\mu}{\partial\beta} + \frac{\partial\mu}{\partial d} \cdot \frac{\partial d}{\partial\beta}$, which can be written as a fraction whose denominator is a squared term and whose numerator is

$$\delta z_C^C + \delta(1 - \mu\lambda)d - \mu(1-\lambda)c - \left(\frac{\delta}{1-\delta}\right)\pi_N e(1\beta(1-\mu\lambda)).$$

This value exceeds

$$\delta z_C^C + (1-\lambda)(\delta d - c) - \left(\frac{\delta}{1-\delta}\right)\pi_N e(1\beta(1-\mu\lambda)).$$

The first two terms are strictly positive (the second is so because d must exceed c/δ) and bounded away from zero for $\pi_C = 1 - \pi_N$ sufficiently large. The third term can be made arbitrarily small by selecting a large enough π_C.

The implication is that lowering β has the effect of lowering μ^*. Because $L^* = \frac{b-\mu^* c}{1-\delta}$, we also obtain that L^* rises. □

References

Saul Alarcon. Personal communication: Wildcoast.

B. Douglas Bernheim and Debraj Ray. Collective dynamic consistency in repeated games. *Games and Economic Behavior*, 1(4):295–326, 1989.

Conservation Stewards Program Conservation International. Conservation agreements: Model, design and implementation. Arlington, Virginia, 2007.

Deacon and Parker. Encumbering...

Engel and Palmer. Payments...

Engel, Pagiola, and Wunder. Designing....

Joseph Farrell and Eric S. Maskin. Renegotiation in repeated games. *Games and Economic Behavior*, 1(4):327–360, 1989.

P. J. Ferraro. Global habitat protection: Limitations of development interventions and a role for conservation performance payments. *Conservation Biology*, 15(4):990–1000, 2001.

P. J. Ferraro and A. Kiss. Direct payments to conserve biodiversity. *Science*, 298:1718–1719, 2002.

Paul J. Ferraro. Asymmetric... a.

Paul J. Ferraro. Cost effectiveness. b.

H. Gjertsen and E. Niesten. Incentive-based approaches in marine conservation: Applications for sea turtles. *Conservation and Society*, 8(1):5–14, 2010.

Hart, Latacz, and Lohmann. Combating...

Latacz, Lohmann, and Van der Hamsvoort. Auctioning....

Raul Lopez. Personal communication: Ejido luis echeverria.

Anne McEnany. Personal communication: International community foundation.

William McShea. Personal communication: Smithsonian institution, October 20 2015.

David A. Miller and Joel Watson. A theory of disagreement in repeated games with renegotiation. *Econometrica*, 2013. Forthcoming.

S. Milne and E. Niesten. Direct payments for biodiversity conservation in developing countries: Practical insights for design and implementation. *Oryx*, 43:530–541, 2009.

John Nash. The bargaining problem. *Econometrica*, 18(2):155–162, April 1950.

Practitioner's Field Guide for Marine Conservation Agreements. The Nature Conservancy and Conservation International, Washington, D.C., final v1 edition, 2009.

E. Niesten and Heidi Gjertsen. Economic incentives for marine conservation. Technical report, Science and Knowledge Division, Conservation International, Arlington, Virginia, 2010.

E. Niesten, A. Bruner, R. Rice, and P. Zurita. Conservation incentive agreements: An introduction and lessons learned to date. Technical report, Conservation International, Washington, D.C., 2008.

Lykhim Ouk. Personal communication: Community engagement manager: Ci-cambodia. Community Engagement Manager.

R. D. Simpson and R. A. Sedjo. Paying for the conservation of endangered ecosystems: A comparison of direct and indirect approaches. *Environment and Development Economics*, 1:241–257, 1996.

K. Svadlenak-Gomez, T. Clements, C. Foley, N. Kazakov, D. Lewis, D. Miguelle, and R.Stenhouse. Paying for results: Wcs experience with direct incentives for conservation. Technical report, TransLinks, Bronx, NY, September 2007.

TNC. Fishery buy-ins: Applying rights-based incentive agreements to sustainable fishery interventions. The Nature Conservancy, Narragansett, Rhode Island, 2013.

S. Troëng and C. Drews. Money talks: Economic aspects of marine turtle use and conservation. Technical report, WWF-International, Gland, Switzerland, 2004.

Joel Watson. Contract and game theory: Basic concepts for settings with finite horizons. *Games*, 4 (3):457–496, September 2013.

Wu and Babcock. Contract....

S. Wunder. The efficiency of payments for environmental services in tropical conservation. *Conservation Biology*, 21(1):48–58, 2004.

S. Wunder. Payments for environmental services and the poor: concepts and preliminary evidence. *Environment and Development Economics*, 13:279–297, 2008.

Ani Youatt. Personal communication: Natural resources defense council.