A contract-theoretic model of conservation agreements

Heidi Gjertsen, Theodore Groves, David A. Miller,
Eduard Niesten, Dale Squires, and Joel Watson*

August 23, 2010

Abstract

We model conservation agreements using *contractual equilibrium*, a concept introduced by Miller and Watson (2010) to model dynamic relationships with renegotiation. The setting takes the form of a repeated principal-agent problem, where the principal must pay to observe a noisy signal of the agent's effort. Lacking a strong external enforcement system, the parties rely on self-enforcement for their relational contract. 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. We argue that the model captures important features of real conservation agreements and reveals the ingredients required for successful agreements.

*Gjertsen and Niesten: Conservation International; Groves, Miller, and Watson: University of California, San Diego; Squires: NOAA Southwest Fisheries Science Center. The authors thank the NOAA Fisheries Service for supporting this research.

Contents

1 Introduction

Incentive-based approaches are increasingly employed by conservation practitioners to encourage changes in resource use (Ferraro 2001; Ferraro and Kiss 2002; Milne and Niesten 2009; Simpson and Sedjo 1996; Troëng and Drews 2004; Wunder 2004, 2008). While past approaches have generally focused on fines and penalties (negative incentives), some current approaches use compensation of various forms (positive incentives) to encourage particular conservation practices. These approaches recognize that conservation can impose a loss in terms of foregone income or access to resources (opportunity cost).

Conservation agreements constitute an important class of positive incentive systems. Under this approach, conservation investors (typically 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 conservation performance. Benefits may be in the form of cash, services, or goods; they are provided periodically upon verification that conservation performance targets are met.

Conservation agreements are of interest because they involve ongoing interaction between the resource users and conservation investors, where the investors provide an ongoing incentive that is external to existing enforcement institutions. The nature of this relationship is fairly unique amongst conservation interventions. Other interventions are either one-shot agreements or 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 in terrestrial settings. For example, Conservation International's Conservation Stewards Program now has sixty four such programs around the world. Applications in marine contexts have received less attention and have yet to receive systematic review, apart from the studies of Niesten and Gjertsen (2010) and Nat (2009). Given that experimentation is fairly recent, there is little empirical evidence to suggest how they perform or how they are best designed and implemented, 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 2008; Hart and Latacz-Lohmann 2004; Latacz-Lohmann and van der Hamsvoort 1997; Wu and Babcock 1996), cost-effectiveness of conservation payments (Ferraro 2002), marine conservation easements (Deacon and Parker 2009), and general conceptual models of payments for environmental services (Engel and Palmer 2008; Engel, Pagiola, and Wunder 2008). However, none of these works directly address the repeated nature of the interactions between the parties to the contract.

This paper develops a contract-theoretic model of conservation agreements and assesses how these agreements can be optimally designed, in particular for marine conservation and fisheries management. A number of case studies are used to guide the development of the model. The layout of the paper is as follows: Section 2 introduces the context through a case study in Cambodia. The modeling approach is outlined in Section 3, followed by analysis and results in Section 4. Section 5 presents additional case studies and Section 6 discusses the implications of our modeling exercise for the design of conservation agreements; Section 6 also describes possible extensions of our model.

2 Case study: Forest protection in Cambodia

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. This is to be achieved through land use zoning and other rules. The parties to the agreement are 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. Alternatively, the parties may transition to a long-term agreement, which would involve reviewing terms periodically but not necessarily every year.

According to the contract, the commune members agreed to not engage in the setting of snares to capture wildlife, to not 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 both 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 of the terms of the agreement, and facilitating effective, transparent and equitable delivery of benefits. Summing up, approximately 8,760 USD is to be spent annually to protect 6555.42 hectares of forest.

Compliance monitoring is to be 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.

During the initial agreement period, a violation of the contract occurred, whereby community members cleared approximately 12 hectares 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 stipulations instead:

1. Before the current agreement can enter into effect, the Council will provide to CI

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%.

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 by two buffalo;
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).

This case illustrates many of the key features of the model developed in the next section. We see that the Cambodia agreement was reached in the context of a dynamic relationship between the NGO and the community. Although the agreement is annual, the parties know that they will continue to interact afterward and they anticipate renewing the agreement; further, the agreement contains provisions for renegotiation each year. Thus, the players view the agreement as an open-ended repeated game. The Cambodia contract is not enforced by an external party such as the government. Rather, it must be self-enforced by the parties to the agreement. That is, compliance must constitute an equilibrium of the repeated game.

Having a workable, enforceable agreement requires a specification of what happens if the agreement is not followed, i.e. penalties for noncompliance. The off-equilibrium penalties must be credible and it must be the case that the parties are better off enforcing the agreement rather than abandoning it. In the model, the community faces a penalty if it violates the agreement and the NGO must make an additional payment if it does not monitor. The community must be induced to protect and the NGO must be induced to monitor. In the Cambodia case, when the community deviates, the NGO reduces the level of benefits but does not eliminate them. In addition, the penalty scheme has been set up such that there is an incentive for community members to report violations by others.

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. The renegotiation made both parties better off than if they had followed the sanction regime specified in the agreement.

3 Modeling approach

Three key features of most conservation agreements are that they are (i) reached in the context of dynamic relationships between non-governmental organizations (NGOs) and local communities, (ii) not formally enforced by any government, and (iii) subject to renewal and renegotiation, where the parties can exercise power. Features (i–ii) suggest modeling conservation agreements using repeated games, in which enforcement power arises endogenously in equilibrium via the threat of adverse changes in how they may play in the future. That is, if a community failed to keep its end of an agreement, the NGO could cease offering any future agreements as punishment, while if the NGO reneged on its promised payments, the community could refuse any future agreements. More generally, we think of conservation agreements as involving two incentive problems, one associated with getting the community to conserve and one associated with getting the NGO to make payments and engage in costly monitoring.

However, feature (iii), the possibility of renegotiation, makes simple threats uncredible. If a conservation agreement was a good idea at the outset of the relationship, then it would still be a good idea at any future date. So if the parties ended up in a punishment phase in which no agreement should be reached, they could jointly deviate to reach a new agreement and make themselves both better off. Doing so, of course, could negate the enforcement power of the punishment phase in the first place. The idea of *contractual equilibrium* (Miller and Watson 2010) is to account for renegotiation within the game, to study endogenous enforcement using threats that are jointly credible. In contrast to previous approaches (various forms of renegotiation proofness, due to Bernheim and Ray 1989, Farrell and Maskin 1989, Pearce 1987, and others), contractual equilibrium explicitly models the process of negotiation and renegotiation. This property gives us the tools to interpret the negotiation outcomes we observe in our case studies.

In a contractual equilibrium, the players meet to negotiate at the start of each period. In any particular period, suppose they know how they would play in the future if they failed to agree, and what their payoffs would be. Then they form an agreement to split the surplus relative to what they would obtain under disagreement. The agreement involves an immediate cash transfer[1] (which is used to split the surplus according to their exogenously given bargaining weights) as well as a recommendation for how they should play going forward. In their agreement, the parties coordinate on strategies that maximize their total surplus, subject to the constraint that their strategies must form an equilibrium in the game from that point on.

If they fail to agree, the players make no immediate transfers and they proceed directly to play their stage-game actions. Since they are under disagreement, in this case there is no particular reason they should be playing efficiently. For example, they may play a Nash equilibrium in the stage game. More generally, the way they play under disagreement must form an equilibrium in the game from that point on, even if their play in the current period is not a Nash equilibrium in the stage game.

Because they negotiate at the beginning of every period, their strategies (under agree-

[1] Equivalently, the transfer may be in-kind, such as water buffalo.

ment or disagreement) include a specification of how they will negotiate in all future periods under all possible future contingencies. Since they always agree in equilibrium, this means that even if they disagree in the current period (this would be a joint deviation that would put them off the equilibrium path), they still anticipate agreeing in the next period.

Miller and Watson (2010) provide an algorithm for constructing a contractual equilibrium using dynamic programming. They show that a contractual equilibrium always exists, and the welfare level (the sum of the players' expected utilities) is the same in all contractual equilibria. In a two-player game, the algorithm constructs the players' agreements in two states: one that is worst for one player, and one that is worst for the other.

In the next section, we explore the implications of contractual equilibrium in a simple model of conservation activity and monitoring.

4 Analysis and results

4.1 The model

The underlying stage game is:

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

where P stands for Protect, E is for "exploit," M is for "monitor," and R is for "rest." Note that monitoring costs the NGO c. We assume that $b \gg e > 0$ and $b \gg c > 0$, so that protection is desirable. This stage game is repeated each period over an infinite horizon. We allow the parties to evaluate their streams of payoffs using different discount factors, δ_C and δ_N, but focus primarily on the case in which $\delta_C = \delta_N = \delta$.

The NGO's action is publicly observed, but the community's action is private. (The NGO does not actually observe its own payoff—it cares about the existence and vitality of the resource regardless of whether it can actually observe the resource. Therefore it cannot infer anything about the Community's action from its own payoff.) There is a monitoring signal $s \in \{G, B\}$ which takes the value B with probability $\lambda > 0$ if $a = EM$, and probability 0 otherwise. That is, a "bad" signal may be observed if the Community exploits the resource and the NGO engages in monitoring, so that monitoring is required to detect exploitation.

Since the community's action is private, the future course of the game must depend only on the observable information, which for each period lies in the set of public signals $S \equiv \{MG, MB, RG\}$. Here, MG represents the outcome in which the NGO monitored and the good signal was realized, MB is the outcome in which the NGO monitored and the bad signal was realized, and RG is the outcome in which the NGO rested and so the signal was inevitably good. Note that RB is not in this set because the bad signal cannot occur when the NGO rests.

4.2 Equilibrium overview

We construct a contractual equilibrium using a strategy with two states. State N is relatively bad for the NGO, while State C is relatively bad for the Community. The *contractual equilibrium value set* (CEV set, V^*) is the line segment spanning the average agreement payoffs in these two states. By randomizing, the players can attain any convex combination of these two states' payoffs. Since under agreement they always optimize, these two states must differ by only a cash transfer, and hence V^* is a line segment with slope -1.

Incentives in both states, under both agreement and disagreement, are provided by specifying which payoff vectors in V^* the players will obtain from the following period as a function of the public signals (we call these the "continuation payoffs"). If the incentives are to be sufficient to enforce the desired behavior, the *payoff span* of V^*—the difference in each player's average payoff between the two states—must be sufficiently large, and the discount factor must be sufficiently high. In this section we sketch the outlines of the equilibrium; the detailed derivation is in Appendix A. The equilibrium is shown graphically in Figure 1.

State *N*, disagreement (Status quo) In State N, under disagreement the NGO rests, the Community exploits, and there is no transfer. Since this is the stage-game Nash equilibrium, no changes in continuation payoffs are needed to enforce it. Hence they simply stay in State N regardless of what happens, so the payoff span does not matter. This is also the status quo before the first agreement is reached, so state N is the initial state.

State C, disagreement (Probation) In State C, under disagreement the NGO monitors, the Community protects, and there is no transfer. This constitutes a kind of "probationary regime," which ends with some probability if either the NGO observes a good signal or the NGO fails to monitor. That is, they stay in State C if MB arises (which can happen only if the Community deviates), and randomize between the states otherwise. In particular, if MG arises (which should happen if neither party deviates), they randomize so as to make the Community e/λ better off than it is in State C; this makes the Community willing to protect. If instead RG arises (indicating that the NGO has deviated), the parties should randomize so as to make the NGO c worse off than if MG had arisen; this makes the NGO willing to monitor. The Community's average discounted utility in state C is exactly e (if it were any lower, the Community would simply deviate to playing E forever); it gets a payoff of zero in the current period and then expects a continuation payoff higher than e. If the payoff span is at least $\frac{1-\delta}{\delta}\left(\frac{e}{\lambda}+c\right)$, then this specification of continuation payoffs is feasible.

Both states, agreement When the parties agree, the Community protects and the NGO mixes, monitoring with probability μ and resting with probability $1-\mu$. Similarly to under disagreement in State C, they structure their continuation payoffs so as to make these actions incentive compatible. If RG arises, they continue in state N to reward the Community, while if MB arises, they continue in state C to punish the Community. If MG arises, they randomize so as to make the NGO's continuation payoff exactly c greater than if RG had arisen; this makes the NGO indifferent between monitoring and not, so it is willing to randomize.

Since monitoring is costly, the parties would like to minimize μ while preserving incentive compatibility. This construction implies that μ is minimized where it sets

$$\frac{1-\delta}{\delta}\left(\frac{e}{\lambda\mu}+(1-\mu)c\right)$$

equal to the payoff span.

The only difference between the two states under agreement is that in state N the NGO makes a much larger transfer—to split the surplus relative to the state N disagreement point of playing ER and staying in state N. In contrast, in state C the disagreement point is much more favorable for the NGO (playing PM and then randomizing between the two states), so the transfer is much smaller.

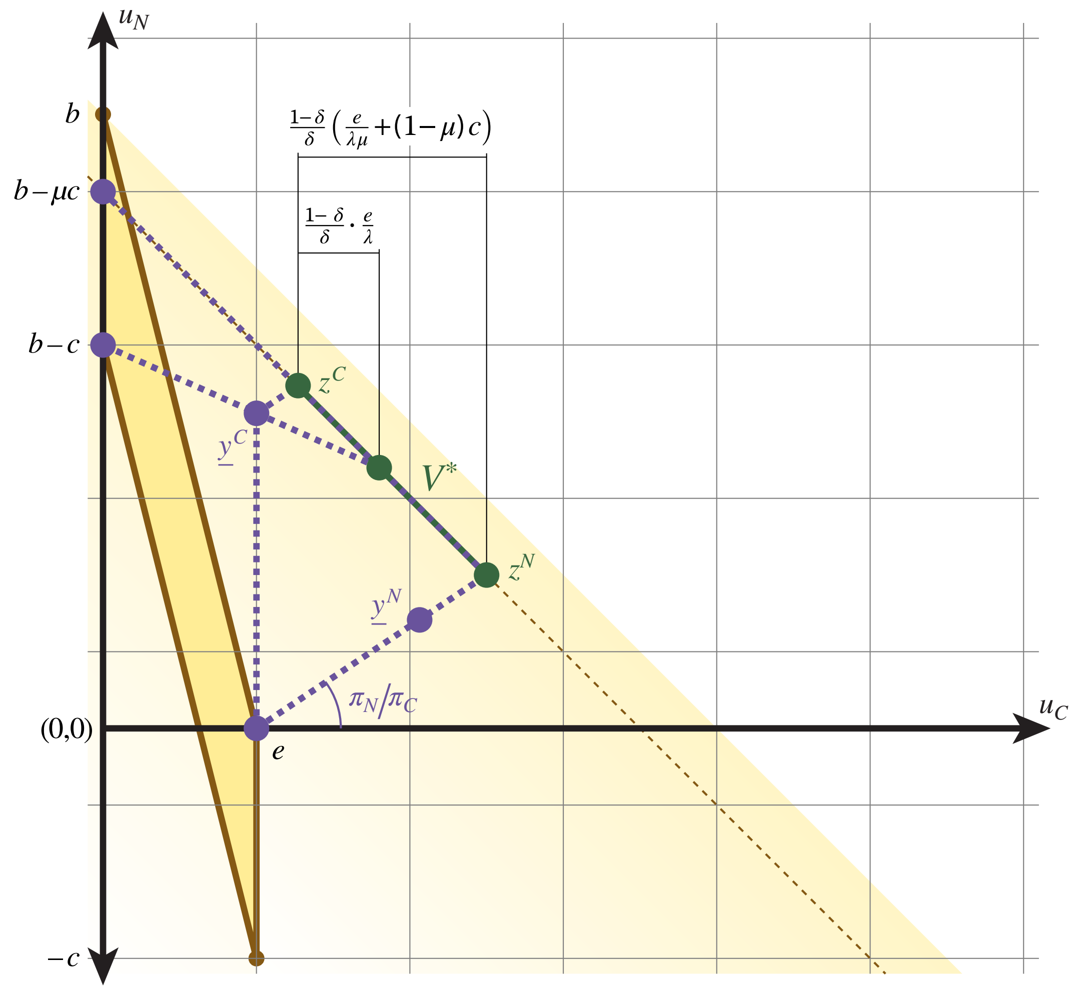

FIGURE 1. CONTRACTUAL EQUILIBRIUM The figure is drawn in terms of both stage game payoff vectors and average discounted utility vectors in the repeated game. The yellow parallelogram outlined in brown is the set of payoff vectors attainable in the stage game without using cash transfers. The lighter yellow triangular region represents the payoff vectors attainable in the stage game *with* cash transfers. 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 $\underline{y}^C$ and $\underline{y}^N$ are attained under disagreement in the two states.

4.3 Implications

Incentives for monitoring When the NGO monitors and the signal is G, it must pay enough to reward the Community for the good outcome. The amount of the payment is increasing in the effort cost and decreasing in the informativeness of the signal. If the NGO were to rest, it would have to pay that amount plus its monitoring cost—otherwise it would not be willing to monitor. This is a key intuition: the NGO's incentives are as relevant as the Community's. If the NGO is not induced to monitor, the Community will not have an incentive to preserve the resource.

Resource protection For the parties to be able to protect the resource in equilibrium, the payoff span must be at least $\frac{1-\delta}{\delta}\left(\frac{e}{\lambda\mu} + (1-\mu)c\right)$. From Figure 1, it is evident that this will be satisfied if δ, π_C, b, and λ are sufficiently high, while π_N, c, and e are sufficiently low. That is, there must be enough surplus to share (δ and b are high while c and e are low), and the Community must obtain enough of the surplus to satisfy its incentives (π_C is high and π_N is low).

Bargaining power and welfare Since the status quo is ER before any agreement is reached, the first agreement starts in state N, and is formed relative to a disagreement point that arises from playing ER. Therefore, as we see from Figure 1, the amount that the NGO pays the Community is larger when the Community has more bargaining power. The same is also true to some extent in state C. However, in state C the disagreement point arises from playing PM, which yields a payoff that is much closer to the Pareto frontier than is the payoff from ER. Hence the effect of the Community's bargaining power is much larger in state N than in state C.

The implication is that V^* will be a longer set when the Community has more bargaining power. What may not be evident from the figure is that when V^* is longer, the community faces a larger punishment for a bad signal, so the NGO doesn't need to monitor as often. Since monitoring is costly, the welfare level is higher when the Community has more bargaining power.[2] By implication, the NGO's payoff is highest when it has an intermediate level of bargaining power—if π_N is too low, the Community appropriates most

[2] If, contrary to our assumptions, the Community's exploitation value exceeded the NGO's benefit of protection net of monitoring cost (that is, if $b < c + e$), then social welfare would instead be increased by shifting bargaining power to the NGO.

of the surplus from protection, while if if π_N is too high, it is either difficult or impossible to motivate the Community to protect.

Heterogeneous discount factors For the case in which the players have heterogeneous discount factors, and are patient enough to achieve protection in equilibrium, it can be shown that μ is decreasing in each δ_i when the δ_j is held fixed (j is the player other than i), but μ is increasing in ε when $\delta_N = \delta+\varepsilon$ and $\delta_C = \delta-\varepsilon$. Similarly, the transfers from the NGO to the Community are increasing in each δ_i, but decreasing in ε. So added patience by either party has beneficial effects both to welfare (by decreasing monitoring costs) and the Community (by increasing the transfers), but a larger difference between the patience of the NGO and the patience of the community can be harmful.

Formal contracts versus contractual equilibrium When the parties agree in a contractual equilibrium, they implicitly agree on the entire contingent future path of their relationship, such as what to do after future disagreements and deviations. However, a contractual equilibrium can be implemented by a sequence of one-period formal contracts, each specifying behavior under agreement in the current period and then behavior under disagreement starting the following period. Once the current formal contract ends, the parties renegotiate (in a way foreseen by their contractual equilibrium) to agree on a new one-period formal contract.

In the contractual equilibrium we construct, a one-period formal contract in state N specifies that if RG is realized then the parties separate at the end of the contract, while if MB is realized then they go to the probationary regime described above. Hence, as long as they stay in state N (no bad signals are realized), they can simply renew their contract as if starting a fresh relationship in the next period.

Renegotiating after a violation If the Community exploits, and the NGO monitors and observes a bad signal, the parties switch to state C. If they were operating under a one-period formal contract, it specifies that they should enter the probationary regime, which is suboptimal because the NGO is supposed to monitor for sure. When they renegotiate, they should reach a new Pareto-improving one-period formal contract in which the NGO pays a small transfer in return for not monitoring so much. Note that this transfer should be much smaller than the transfer that is paid in state N, since the probationary regime puts the Community in a much weaker bargaining position.

Bad signals Given the monitoring structure, as long as the Community protects the resource, a bad signal will never arise, and they will stay in state N. We don't want to draw any implication from this property, since this particular monitoring structure is a bit extreme. In a more general model, bad signals could arise with some baseline probability even when the NGO rests, and, indeed, even when the Community protects. However, information from the case studies reported here indicates that, when there is preliminary evidence of a bad signal, the parties typically engage in a round of additional fact-finding to determine whether the Community really exploited, implying that false negatives are rare.

5 Additional case studies

In this section we present three additional cases of conservation contracts applied to the marine context. The cases were researched as part of a project for Conservation International's Marine Management Area Science Program and are presented in greater detail in Niesten and Gjertsen (2010). We discuss how these cases conform or differ from the assumptions and results of the model.

5.1 Marine protection in Fiji

Navini is one of 32 small islands in Fiji's Mamanuca group.[3] This region has been the focus of tourism development since the mid-1960s, and the Navini Island Resort, established in 1976, is one of the smallest tourism establishments. Navini Island's coral reefs have suffered various threats including over-fishing, destructive fishing through the use of explosives and poison, and pollution. Communities depended heavily on marine resources for income and protein. Uncontrolled fishing by locals and outsiders led to decreases in fish stocks and degradation of the ecosystem. Pressure was due mainly to limited livelihood options in the area and increasing cash needs for expenses such as school fees. This pressure has been reduced more recently as the majority of households in nearby communities now depend heavily on the tourism industry for their livelihoods.

Situated in Fiji's busiest tourism area and close to several fisheries landing sites, the Navini Island Resort management team values the reef system surrounding the island and

[3]Descriptive information in this section is summarized from a report submitted to Conservation International by Patrick Saki Fong.

the need for conservation to sustain tourism development. In 1988, the resort management team approached the chief and members of the Tui Lawa clan to ask for protection of the reefs surrounding the island. The team from Navini explained the rationale for establishing the MPA and discussed the commercial use and value of fish being taken from the area. The resort then offered payment in excess of that value as compensation for establishing the MPA, which resulted in the agreement to establish what is now Fiji's oldest MPA. The Navini Island marine protected area is thus a conservation agreement between the Navini Island Resort and the chiefly clan of the Tui Lawa.

The agreement and MPA are not formally recognized under Fijian government legislation or policies. The reefs surrounding the island resort form the MPA and are fully protected under a complete ban on extraction of any marine resources. The agreement is renewable in every 12 months. Each year the management of the resort holds a meeting with the Tui Lawa and decides on the renewal and the terms of agreement for the following year. The resort pays approximately 215 USD annually for the renewal of the agreement. These funds are managed by the Tui Lawa and used for education and development needs of clan members. The resort also provides library books and other school materials for the local primary school and supports community development projects through financial or technical assistance. The total value of these benefits exceeds 400 USD per year.

Since the MPA is close to the island resort, surveillance and monitoring of illegal activities falls under the daily duties of Navini resort staff, though there is no formal surveillance plan. When poachers are detected, the staff explains that the area is a no take zone and requests that they leave. Once the poachers depart, the staff informs the Tui Lawa, who uses his traditional authority to speak to the poachers or the chief responsible for the poachers. Between 1989 and 1992, after the establishment of the MPA and recovery of marine resources, some poaching took place by village fishermen. The Tui Lawa was alerted to this problem, and poaching diminished soon thereafter. Poaching incidents are now rare, and usually do not involve locals.

The Navini contract is between a community and a resort company, which we will consider as the "NGO" from our model. The lease period is one year, with renegotiation happening after each period. These features are consistent with the repeated-game context of the model. The model assumes that there is no external enforcement. This is appropriate for Navini where external government enforcement is simply not available. Instead, the resort and the clan brokered an agreement that is expected to be self-enforced. The form of monitoring specified in the model is likely to be appropriate for this case, with some probability of a bad signal being observed if the community exploits (i.e. the resort staff may detect exploitation through their low-level informal monitoring). As in the model, it is unlikely that a false positive will occur through this form of monitoring, unless there is potential for community members to be conducting activities in the MPA that may be mistaken for poaching. One difference with the model is that since the monitoring choice is not a binary one—i.e. the NGO can choose a level of monitoring that is between "monitor" and "rest"—monitoring can be enforced by implicit threats off the equilibrium path, rather than occasional payments on the equilibrium path. As long as the NGO continues to monitor, these threats need not be carried out.

In the model, reaching an agreement requires that it generate enough surplus for the parties to share, and the community must obtain a sufficient share. Many conservation agreements assess whether there is an opportunity for gains from trade by measuring opportunity cost as a basis for payment, with donor willingness to pay as an upper bound. For example, in Navini, the value of the fish catch from the area to be protected was used as an estimate of the opportunity costs to the clan of declaring a marine protected area. This informed the size of the payment needed for the community to participate in the agreement. Both the Navini Island Resort and the clan exercised bargaining power in this case. It is likely that both parties are relatively patient (i.e., have relatively high discount factors). The resort has invested in the area and its success depends on maintaining successful protection. The clan has chosen to receive benefits in the form of education and community development, suggesting a long-term view.

5.2 Turtle preservation in the Solomon Islands

Tetepare Island is located in Western Province of the Solomon Islands and is approximately 11,880 hectares in size. The landowners—collectively the Tetepare Descendants' Association (TDA)—live throughout the Solomon Islands, but the majority live in four villages on Rendova, the largest (approximately 40,000 hectares) and closest island west of Tetepare. In 2002, Australian biologists and a TDA project officer initiated a conservation agreement to protect leatherback turtles. The program involves a series of payments related to turtle nesting and hatching that is undisturbed by community members. The participating villages each select a turtle monitor that is responsible for recording data about turtle nesting activities and receive financial incentives each time they do so. The monitors do not receive a salary.

The program operated as follows in 2004: A villager who observes a leatherback nesting reports it to the turtle monitor. If the monitor tags the turtle and records the information on the data sheet, the observer is paid 2 USD and the monitor is paid 1.33 USD. A lower payment is made to a villager that finds a nest or tracks after the turtle has gone back to sea. If the observer disturbs the turtle in any way she does not receive the payment. In addition to recording the data, the monitor photographs the turtle, recording the date and time for verification purposes. In addition to the individual payments, 1.33 USD is placed in a community fund. If the nest successfully hatches (i.e., at least one hatchling emerges), the initial reporter is paid 4 USD, the monitor 1.33 USD, and an additional 4 USD is paid into the community fund. The fund is managed by a Board of community members. There are five signatories to the fund, all of whom must sign to withdraw money. When funds are withdrawn, the signatories are required to provide minutes of the community meeting clarifying how the money will be used. The community fund has only recently been opened, and the community has indicated that they are saving up to build a new community hall and/or school classroom.

A TDA staff member makes six visits to the villages per season and collects the photos and data sheets and distributes the payments. According to the local project coordinator, there is great competition to find and report the turtles and nests. This project creates multiple incentives for turtle conservation. Every villager and turtle monitor has the potential to access payments for reporting and not disturbing or consuming turtles or their eggs. In addition, there is an incentive to prevent others from disturbing the turtle or the nest, since they receive an additional payment if the nest hatches. Finally, the community as a whole faces an incentive in the form of contributions to the fund if turtles or their nests remain undisturbed. Thus, someone who did not find the nest receives some benefit (via the community fund) from not harvesting. This of course depends on the degree to which everyone perceives the fund as something from which they will benefit.

Since 2007, tides have been quite high, inundating many of the nests. Therefore, all nests below the high water line are now relocated to a safe zone that is monitored by the turtle monitors. The finders still receive a payment, but do not receive a payment for hatching. Instead, the hatching payments all accrue to the community fund. TDA has recently requested an increase in the level of payments because of inflation.

The conservation agreement has been successful in protecting nesting leatherback females and their eggs. While there are occasional poaching incidents, previous to the project nearly all eggs and females were harvested. In addition, the project has provided modest income to villagers and funds for community development. Current efforts to secure long-term funding in the form of an endowment are showing some promise. In addition to financial incentives linked to turtle monitoring, the project also provides incentives in the form of a scholarship program and conservation and tourism-based employment.

The attempt by Rendova to secure long-term financing through an endowment provides evidence of a repeated long-term game. The agreement has been upheld without interruption for eight years. In this case, violations on the part of the community involve a simple reduction in benefits, as the individual and community receives no payment for any nest or turtle that is harvested. The model assumes that there is no external enforcement. This is appropriate for Rendova, where external enforcement is not available. In Rendova the agreement has provided additional incentives for self-enforcement, as individuals prevent others from poaching eggs for which they will obtain a future payment. The monitoring in this agreement is similar to the model. The monitors can choose to respond to a report by villagers or not, but in the case that they do not, they forego a payment. Thus, the agreement provides incentives to monitor. In the model, the less the parties discount the future, the less they must be penalized in the short-term, because of the value of the long-term relationship. Of the cases in this paper, Rendova is likely to exhibit the shortest time horizon. Thus, the payments must be large enough to induce the community to protect. The size of a single payment for turtle eggs and nests relative to the benefits from exploitation is likely to be quite large, as exploitation is for subsistence purposes and there are alternative protein sources available and the payment compares favorably to the daily wage rate. Despite the large payments required by the community, there were still gains from trade to be captured in an agreement since the benefits to the NGO from protection are so large. In other words, the willingness to pay for protection is significantly greater than the opportunity cost.

5.3 Grey whale habitat protection in Mexico

Laguna San Ignacio is situated on the Pacific Coast of Baja California Sur, Mexico. The 80,000 hectare lagoon forms the southern boundary of the Vizcaíno Desert. Laguna San Ignacio 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. Thousands of gray whales make an annual 10,000-mile voyage from feeding grounds in the Arctic circle to the warmer waters of Laguna San Ignacio to calve and rear their young before journeying back to Alaska in the spring. In addition to providing grey whale habitat, Laguna San Ignacio hosts at least 221 species, including numerous birds, green sea turtles, and bottlenose dolphins.

In 1994, Mitsubishi proposed to establish a salt plant at Laguna San Ignacio. The proposal was eventually defeated in 2000 through the efforts of local and international NGOs, but coastal development pressure continued to threaten the area. To conserve the area over the long term, Mexican NGO Pronatura suggested the option of an easement, which would come with the enforcement power of Pronatura lawyers. 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 Echeverría Alvarez on the southern shore of Laguna San Ignacio.

There are four parties to the agreement, each with a specific role. Ejido Luis Echeverría agrees to limit coastal development. Pronatura monitors compliance. The International Community Foundation (ICF) is a San Diego foundation responsible for disbursing funds to Ejido Luis Echeverría. They maintain a trust fund and manage it as a third party so there is transparency and accountability. Maijanu is an NGO that was created in Ejido Luis Echeverría 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 area and survey the terrain and interview people. If Pronatura determines that the Ejido has met its obligations, ICF will disburse to Maijanu the annual interest generated from the Ejido Luis Echeverría Alvarez Seed Fund, which is capitalized in the amount of 650,000 USD. These annual payments amount to approximately 25,000 USD per year. The Ejido 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 Ejido leadership and all the members vote in a general assembly for the proposals.

If the Ejido's obligations in the contract are not met, the payments to the Ejido will not be disbursed. If the violation created damage that can be restored, the payments can begin again when 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. ICF maintains a Legal Defense Fund of 225,000 USD to enforce and defend the terms and conditions of the conservation easement. ICF also maintains a Stewardship Fund of 250,000 USD that disburses 10,000 USD per year to Pronatura for monitoring. Thus far, the terms of the easement have been met by Ejido Luis Echeverría, and funds have been released to the Ejido for community projects, including technical assistance for raising chicken and goats, improvements to the Ejido store, and pilot projects for cactus and artisanal salt production.

Intense development pressure in Baja California means that conservation requires incentives to compete with potentially lucrative alternative land use. The Ejido Luis Echeverría conservation easement recognized these opportunity costs of forgoing development and created incentives for the Ejido to protect valuable whale habitat in perpetuity. Responsibilities of the various stakeholders and the associated procedures (monitoring, reporting, enforcement, payments, etc.) are clearly spelled out in a written contract. Long-term funding was secured up-front in trust funds (managed by a third party) so that all easement-related costs can be covered in perpetuity. The various required activities and contingencies were all taken into account in the design. Because payments are only released when conservation objectives are met (based on monitoring compliance with the terms of the easement), incentives are performance-based and these costs are not incurred if conservation is not achieved. Furthermore, a legal team and funds are available in case infractions must be prosecuted.

In terms of the benefits to the Ejido from engaging in conservation, there may be more pressure on the arrangement in the future if land speculation raises the opportunity cost of conserving the land. The agreement does not currently include provisions to increase the annual benefits to the Ejido. Finally, the easement does not provide direct benefits to non-Ejido members. Land disputes and tension between these residents and Ejido members may threaten the success of the easement if, for example, illegal activities by non-members result in negative monitoring results that prevent payments. However, this has not occurred to date, and given the non-Ejido members' limited political and economic status it may not be in their best interest to interfere with the easement.

Pronatura has attempted to implement these methods in other Laguna San Ignacio ejidos, but encountered a number of challenges. Many areas expect larger future benefits from developing the land (because of speculation), therefore the land is overvalued making negotiations difficult and expensive. The legal status of some land is unclear and there are a number of absentee landlords, again complicating the negotiations. Other areas would probably require (large) individual payments from an easement, as funds for small-scale community development projects are unlikely to be sufficiently attractive. Conservationists are considering other approaches for these areas, since easements may be unfeasible or prohibitively costly.

In Laguna San Ignacio, the agreement is made in perpetuity and the NGO secured long-term funding (in the form of trust funds) for all components of the agreement, including monitoring and community payments. The agreement does not contain provisions for renegotiation; however, one might expect renegotiation to occur if the community violates the agreement and it is costly for the NGO to punish them in the courts according to the terms of the agreement. Although this case is more complex, with four parties to the agreement, the basic setup of a community that chooses to exploit or protect and an NGO that monitors or rests applies.

One specification of the model that does not apply to the Laguna San Ignacio case is the assumption of no external enforcement, as the agreement contains provisions and funds to settle disputes in court. If external enforcement is frictionless and evidence is available, then there is no need for self-enforcement. If external enforcement is costly, then parties would settle (renegotiate).

The monitoring specified in the model is likely to be appropriate for this case, with some probability of a bad signal being observed if the community exploits. As in the model, it is unlikely that a false positive will occur, since the NGO is not likely to observe coastal development if it has not occurred. Similar to the case of Navini, the NGO is able to choose a level of monitoring that is between "monitor" and "rest," thus avoiding the need for a negative incentive of additional payments to the community if it rests. Of the cases in this paper, Laguna San Ignacio is likely to exhibit the longest time horizon. In Laguna San Ignacio, the payments that are made to the community are quite low relative to the outside option of selling the land to developers. If the community had a shorter time horizon they would require a much higher payment. However, if the discount rate is low enough, 25,000 USD per year can compete with the money they could earn now by selling their land.

6 Conclusion

This paper developed a contractual equilibrium model of conservation agreements. The model captures key characteristics of real-world conservation agreements, in particular: (i) ongoing relationships between non-governmental organizations (NGOs) and local communities, (ii) lack of external enforcement, and (iii) renegotiation. Conservation agreements are modeled using repeated games with renegotiation. Notably, we capture two incentive problems inherent in conservation agreements: one associated with inducing the community to conserve, and the other to induce the NGO to engage in costly monitoring.

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 (some control over the renegotiation process).

3. The NGO and community must have high enough discount factors to achieve protection in equilibrium. While added patience by either party is beneficial, a larger difference between the two discount factors can be problematic.

4. Parties should anticipate how their agreements will be renegotiated over time, in particular following any infraction.

These implications of the model suggest how we might structure future theoretical and empirical work. In particular, we plan to select and classify case studies that vary along the following dimensions:

1. Incentives to monitor and the nature of monitoring that is required,

2. Determinants of bargaining power, especially in terms of who has control over the renegotiation process, and

3. Discount factors of the two parties, in terms of both levels of patience and the viability of a long-term relationship (e.g. the probability that funding will continue to be made available to the NGO).

Finally, we mention some extensions of our model. We regard the last two as worthwhile topics for future work, which we have begun to explore.

1. Bad signals can arise even when the Community protects or the NGO rests

   The monitoring structure in our current model is special in the sense that bad signals can arise only when the parties play EM. This makes the equilibrium construction simple and transparent. However, if bad signals can arise even when the Community protects, then the Community will sometimes have to punished even when it is innocent. This realistic feature is widely recognized in the literature on repeated games with imperfect monitoring.

2. The NGO has outside options

   The current model assumes that the NGO can reach a conservation agreement with only this particular Community. However, if there are many relevant communities and the NGO does not have the funds to reach agreements with all of them, then the NGO's threat to walk away from any particular agreement is credible. This credible threat puts a lower bound on the NGO's payoff, and therefore an upper bound on the Community's payoff. Such a bound is likely to bind only when the Community already faces strong incentives, so it is likely to be beneficial to the NGO.

3. The resource is a stock with growth-depletion dynamics

   The current model does not captures the kinds of dynamics that are inherent in resource protection problems. If the resource is non-renewable, a deviation by the Community may permanently reduce the surplus available in the relationship. Even if the resource is renewable, it may take a long time to recover from partial exploitation. Discouraging exploitation of non-renewable or gradually-renewable resources may require higher-powered incentives.

4. The Community is a group of players with collective-action problems

   In most settings, the Community consists of many individuals, each of whom must be motivated to protect the resource. If the Community as a whole faces strong enough incentives, it can motivate its members through the threat of social sanctions. If these sanctions are expensive, however, the NGO can help motivate community members through individual-level rewards and punishments, as in the case study of forest protection in Cambodia.

A Detailed construction

Here we provide the construction of the contractual equilibrium value (CEV) set for the general case in which the players may have different discount factors. Let $A \equiv \{PM, PR, EM, ER\}$ denote the set of stage-game action profiles. For each player $i \in \{C, N\}$, let $u_i : A \to \mathbb{R}$ denote the stage-game payoff function. That is, $u_C(PM) = 0$, $u_N(PM) = b - c$, and so on. We extend u_i to be defined over mixed actions.

Since the players have different discount factors, we express their payoffs in total terms (that is, present discounted value), rather than average terms, in this appendix. Quantities expressed in total terms are written with a tilde to distinguish them form quantities expressed in average terms in the text. For the special case in which both players share the same discount factor δ, the results in the text arise from multiplying total payoffs by $1 - \delta$. From a given period of time, player i's continuation payoff can be written

$$m_i + u_i(a) + \delta_i \mathbb{E}\left[\tilde{g}_i(s)|a\right],$$

where m_i is the monetary transfer made to player i in the negotiation phase of the current period, $a \in A$ is the action profile played in the current period, $s \in S$ is the public signal realized in the current period, and $\tilde{g}_i(s)$ is i's continuation value from the start of the next period in total discounted terms. Monetary transfers between the players are balanced, so that $m_C + m_N = 0$. The function $\tilde{g}_i$ represents how the players coordinate their behavior in future periods as a function of the outcome of current-period interaction.

Because of the availability of immediate—and therefore undiscounted—monetary transfers at the time of negotiation in each period, the CEV set is a line segment with slope -1. It is thus written, in total utility terms, as $\tilde{V}^* = \text{co}\{\tilde{z}^C, \tilde{z}^N\}$ where "co" denotes "convex hull," $\tilde{z}^N$ is the endpoint that favors the Community (which arises in state N under agreement) and $\tilde{z}^C$ is the endpoint that favors the NGO (arising in state C under agreement). Because the slope is -1, we have $\tilde{z}^C_C + \tilde{z}^C_N = \tilde{z}^N_C + \tilde{z}^N_N$. The *level* of the CEV set is

$$\tilde{L}^* \equiv \tilde{z}^C_C + \tilde{z}^C_N = \tilde{z}^N_C + \tilde{z}^N_N$$

and the *payoff span* is

$$\tilde{d}^* \equiv \tilde{z}^N_C - \tilde{z}^C_C.$$

We can find these extreme values by using the Miller-Watson algorithm described below.

In every period of a contractual equilibrium (regardless of what occurred previously), renegotiation ensures that the players will achieve the endogenously determined joint value $\tilde{L}^*$ regardless of their history of interaction. Note, however, that the history will influence the *division* of this value between the players by the way that it determines the *disagreement point* of their negotiation. Thus, it is generally the case that $\tilde{g}(s) \neq \tilde{g}(s')$, but we have $\tilde{g}_C(s) + \tilde{g}_N(s) = \tilde{L}^*$ for all $s \in S$. Subtracting this constant as well as the transfer m made earlier in the current period, the players' expected payoffs starting in the action phase can be expressed as $\tilde{w}(\alpha, \tilde{\eta}) \equiv (\tilde{w}_C(\alpha, \tilde{\eta}), \tilde{w}_N(\alpha, \tilde{\eta}))$, where $\alpha \in \Delta A$ is a mixed stage game action profile, and $\tilde{\eta} : S \to \mathbb{R}$ describes the value taken from the NGO and given to the Community, starting from the next period, as a function of the realized signal:

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

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

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

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 $d \geq 0$:

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

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 ($\tilde{\eta}^N(s) = 0$ for all s), solves (1). The maximized value is $\tilde{\gamma}^N(\tilde{d}) = \frac{\pi_N}{1-\psi}e$.*

Proof. Because ER is a Nash equilibrium of the stage game, α^N and $\tilde{\eta}^N$ satisfy the conditions 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:

$$(-c) + \delta_N \left(\lambda(-\tilde{\eta}(MB)) + (1-\lambda)(-\tilde{\eta}(MG))\right) \geq 0 + \delta_N(-\tilde{\eta}(RG)). \tag{2}$$

The value of the objective function is thus $\frac{1}{1-\psi}\big(\pi_N e - \mu\pi_C(-c) + \psi\mathbb{E}\left[\tilde{\eta}(s)|E,\mu\right]\big)$. Note that $\mathbb{E}\left[\tilde{\eta}(s)|E,\mu\right] = \mu\lambda\tilde{\eta}(MB) + \mu(1-\lambda)\tilde{\eta}(MG) + (1-\mu)\tilde{\eta}(RG)$. Since $\tilde{\eta}(s) \in \left[-\tilde{d}, 0\right]$ is required, Eq. 2 implies that $\lambda\tilde{\eta}(MB) + (1-\lambda)\tilde{\eta}(MG) \leq -\frac{c}{\delta_N}$, which further implies that $\mu\lambda\tilde{\eta}(MB) + \mu(1-\lambda)\tilde{\eta}(MG) + (1-\mu)\tilde{\eta}(RG) \leq -\frac{\mu c}{\delta_N}$. 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 $\tilde{\eta}^N$ solve the maximization problem that defines $\tilde{\gamma}^N$, and the resulting value of $\tilde{\gamma}^N$ 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 $d \geq 0$:

$$\begin{aligned} \tilde{\gamma}^C(d) \equiv \min_{\tilde{\eta},\alpha} & \frac{\pi_N u_C(\alpha) - \pi_C u_N(\alpha) + \psi\mathbb{E}\left[\tilde{\eta}(s)|\alpha\right]}{1-\psi}, \\ \text{s.t.} & \begin{cases} \tilde{\eta} : S \to \left[0, \tilde{d}\right], \\ \alpha \in \Delta^{\mathrm{U}}A \text{ is a Nash equilibrium of } \langle A, \tilde{w}(\cdot,\tilde{\eta})\rangle. \end{cases} \end{aligned} \tag{3}$$

Lemma 2. *Suppose that*

$$\tilde{d} \geq \frac{c}{\delta_N} + \frac{e}{\delta_C\lambda}. \tag{4}$$

Then the stage game action profile $\alpha^C = PM$, combined with $\tilde{\eta}^C(MB) = 0$, $\tilde{\eta}^C(MG) = \frac{e}{\delta_C\lambda}$, and $\tilde{\eta}^C(RG) = \frac{c}{\delta_N} + \frac{e}{\delta_C\lambda}$, solves (3). The minimized value is

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

On the other hand, if (4) does not hold then $\tilde{\gamma}^C(d) = \frac{\pi_N}{1-\psi}e = \tilde{\gamma}^N(d)$.

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 $(b-c) + \delta_N(-\tilde{\eta}(MG)) \geq b + \delta_N(-\tilde{\eta}(RG))$, which simplifies to

$$\tilde{\eta}(RG) \geq \tilde{\eta}(MG) + \frac{c}{\delta_N}. \tag{5}$$

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

$$\tilde{\eta}(MG) \geq \tilde{\eta}(MB) + \frac{e}{\delta_C\mu\lambda}. \tag{6}$$

With the specified stage-game action profile, the value of the objective function is $\frac{1}{1-\psi}\big(-\pi_C(b-\mu c) + \psi\mathbb{E}\left[\tilde{\eta}(s)|P,\mu\right]\big)$. Note that $\mathbb{E}\left[\tilde{\eta}(s)|P,\mu\right] = \mu\tilde{\eta}(MG) + (1-\mu)\tilde{\eta}(RG)$. Since this value is increasing in $\tilde{\eta}(MG)$ and $\tilde{\eta}(RG)$, and since we have a minimization problem with $\tilde{\eta}(s) \in \left[0, \tilde{d}\right]$, it is optimal to have Eq. 5 and 6 hold with equality and to have $\tilde{\eta}(MB) = 0$. Thus $\tilde{\eta}(MB) = 0$, $\tilde{\eta}(MG) = \frac{e}{\delta_C\mu\lambda}$, and $\tilde{\eta}(RG) = \frac{c}{\delta_N} + \frac{e}{\delta_C\mu\lambda}$. Plugging these values into the objective function, we obtain the value

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

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 $\tilde{\eta}(MB)$, $\tilde{\eta}(MG)$, and $\tilde{\eta}(RG)$ to be feasible. □

The next step in the Miller-Watson algorithm is to calculate the maximal fixed point of the function $\tilde{\Gamma} \equiv \tilde{\gamma}^N - \tilde{\gamma}^C$. From Lemmas 1 and 2, we see that

$$\tilde{\Gamma}(\tilde{d}) = \frac{1}{1-\psi} \cdot \begin{cases} \pi_N e + \pi_C(b-c) - \psi\frac{e}{\delta_C\lambda} & \text{if (4) holds} \\ 0 & \text{otherwise.} \end{cases}$$

Rearranging terms and simplifying, we see that the first line is greater or equal to the second if and only if

$$\pi_N e + \pi_C b \geq (\pi_C \delta_N + \pi_N \delta_C) \frac{e}{\delta_C \lambda} + \pi_C c. \tag{7}$$

We can then write the maximal fixed point of $\tilde{\Gamma}$ as:

$$\tilde{d}^* = \frac{1}{1-\psi} \cdot \begin{cases} \pi_N e + \pi_C (b-c) - \psi \frac{e}{\delta_C \lambda} & \text{if Eq. 7 holds} \\ 0 & \text{otherwise.} \end{cases}$$

This number $\tilde{d}^*$ is the payoff span of the CEV set.

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

$$\begin{aligned} \tilde{\rho}(\tilde{d}^*) \equiv \max_{\tilde{\eta},\alpha} \; & u_C(\alpha) + u_N(\alpha) + (\delta_C - \delta_N)\mathbb{E}_s \left[\tilde{\eta}(s) + (\delta_C - \delta_N)\tilde{d}^* \middle| \alpha \right], \\ \text{s.t.} \; & \begin{cases} \tilde{\eta} : S \to [0, \tilde{d}^*], \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \langle A, \tilde{w}(\cdot, \tilde{\eta}) \rangle. \end{cases} \end{aligned} \tag{8}$$

It is not difficult to show that if $|\delta_C - \delta_N|$ is sufficiently small, then Problem 8 is solved by having the Community select P with probability one. We assume that this is the case. Since we know that the stage-game optimum PR cannot be enforced by any $\tilde{\eta}$, we conclude that, in order for the Community to play P, the NGO must play M with some probability $\mu > 0$. So Eq. 8 becomes:

$$\begin{aligned} \tilde{\rho}(\tilde{d}^*) \equiv \max_{\tilde{\eta},\alpha} \; & b - \mu c + (\delta_C - \delta_N)\big(\mu \tilde{\eta}(MG) + (1-\mu)\tilde{\eta}(RG)\big) + (\delta_C - \delta_N)\tilde{d}^*, \\ \text{s.t.} \; & \begin{cases} \tilde{\eta} : S \to [0, \tilde{d}^*], \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \langle A, \tilde{w}(\cdot, \tilde{\eta}) \rangle. \end{cases} \end{aligned} \tag{9}$$

Recalling the analysis from the proof of Lemma 2, we see that Eq. 5 is required for the NGO to choose M with positive probability. Furthermore, Eq. 6 is required for the Community to select P. It is optimal to select μ as low as possible because it serves to increase $b - \mu c$ and, assuming that δ_C and δ_N are close, this dominates the change in

$\mu\tilde{\eta}(MG) + (1-\mu)\tilde{\eta}(RG)$. We thus have that Eq. 5 and 6 hold with equality, and we also have the condition that $\tilde{\eta} \in [0, \tilde{d}^*]$. The lowest value of μ consistent with these conditions is that μ^* which solves $\frac{c}{\delta_N} + \frac{e}{\delta_C \mu^* \lambda} = \tilde{d}^*$, which simplifies to

$$\mu^* = \frac{\delta_N e}{\lambda \delta_C \big(\delta_N \tilde{d}^* - c\big)}. \tag{10}$$

The Miller-Watson algorithm identifies the endpoints of the CEV set in terms of $\tilde{d}^*$, using the functions $\tilde{\gamma}^C$ and $\tilde{\rho}$. The endpoint favoring the NGO is given by:

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

We then have $\tilde{z}^{N*} = \tilde{z}^{C*} + (\tilde{d}^*, -\tilde{d}^*)$ and $\tilde{L}^* = \tilde{z}_C^{C*} + \tilde{z}_N^{C*} = \tilde{z}_C^{N*} + \tilde{z}_N^{N*}$. The expressions for these quantities in terms of parameters are quite complicated, except in the special case of equal discount factors.

The case of equal discount factors In the special case of equal discount factors, i.e., $\delta_C = \delta_N \equiv \delta$, Eq. 7 becomes

$$\pi_N e + \pi_C(b-c) - \frac{e}{\lambda} \geq (1-\delta)(\pi_C b - \pi_N c). \tag{11}$$

The payoff span of the CEV set is given in average terms as:

$$d^* = \begin{cases} \pi_C(b-c) + \pi_N e - \frac{e}{\lambda} & \text{if (11) holds,} \\ 0 & \text{otherwise.} \end{cases}$$

Assume that Eq. 11 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 $\gamma^C = \pi_N e + \pi_C(b-c) - \frac{e}{\lambda}$.

Then the CEV endpoint vectors are

$$z^{C*} = \pi\rho(d^*) + (1 - \delta\big(-\gamma^C(d^*), \gamma^C(d^*)\big),$$
$$z^{N*} = \pi\rho(d^*) + \big(d^* - \gamma^C(d^*), \gamma^C(d^*) - d^*\big),$$

where $\pi = (\pi_C, \pi_N)$.

References

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

R. T. Deacon and D. P. Parker. Encumbering harvest rights to protect marine environments: A model of marine conservation easements. *Australian Journal of Agricultural and Resource Economics*, 53:37–58, 2009.

S. Engel and C. Palmer. Payments for environmental services as an alternative to logging under weak property rights: The case of Indonesia. *Ecological Economics*, 65:799–809, 2008.

S. Engel, S. Pagiola, and S. Wunder. Designing payments for environmental services in theory and practice: An overview of the issues. *Ecological Economics*, 65:663–674, 2008.

Joseph Farrell and Eric 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. The cost-effectiveness of conservation payments. *Land Economics*, 78(3):339–353, 2002.

P. J. Ferraro. Asymmetric information and contract design for payments for environmental services. *Ecological Economics*, 65:810–821, 2008.

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

R. Hart and U. Latacz-Lohmann. Combating moral hazard in agri-environmental schemes—A multiple-agent approach. *European Review of Agricultural Economics*, 32(1):75–91, 2004.

U. Latacz-Lohmann and C. P. C. M. van der Hamsvoort. Auctioning conservation contracts: A theoretical analysis and an application. *American Journal of Agricultural Economics*, 79:407–418, 1997.

David A. Miller and Joel Watson. A theory of disagreement in repeated games with renegotiation. Working paper, January 2010.

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

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.

David Pearce. Renegotiation-proof equilibria: Collective rationality and intertemporal cooperation. Working paper, December 1987.

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.

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

J. Wu and B. A. Babcock. Contract design for the purchase of environmental goods from agriculture. *American Journal of Agricultural Economics*, 78:935–945, 1996.

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.