Title
A Contract-theoretic Model of Conservation Agreements

ISBN
9789814513463

Authors
Gjertsen, Heidi
Groves, Theodore
Miller, David A
et al.

Publication Date
2014-06-01

DOI
10.1142/9789814513470_0015

Conservation Agreements: Relational Contracts With Endogenous 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

Revised December 2019

In press, *Journal of Law, Economics, and Organization* (37, 2021)

Abstract

This paper examines the structure and performance of *conservation agreements*, which are relational contracts 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. A repeated-game model is developed and utilized to organize an evaluation of real conservation agreements, using three case studies as representative examples.

*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 (Hadri and Mikhail 2014). Groves is especially grateful to Professor William Mikhail for his many courtesies and gracious hospitality. The authors also appreciate helpful comments from Bård Harstad, Lisa Bernstein, seminar participants at UC Berkeley, and conference participants at the 2016 University of Oslo Theory and Environment Workshop. The authors are most grateful to the editor and three anonymous referees for their thoughtful guidance and detailed comments. Miller and Watson thank Isla Globus-Harris, Jake Johnson, and Erik Lillethun for expert research assistance.

Contents

1 Introduction

Most long-term contractual relationships (such as partnership, employment, and buyer-supplier relationships) rely on a measure of self-enforcement. The growing literature on *relational contracts* builds on the framework of repeated games to study the basic incentive problems that these relationships face. Technical advances in this literature include incorporating transfers and negotiation into a multi-stage account of interaction within each period of time.[1] Prominent applications in the relational-contracting literature include employment relationships and commercial supplier relationships.

We apply relational-contract theory to *conservation agreements*, an important class of incentive systems for protecting natural resources. A typical conservation agreement involves an environmental resource, such as a rain forest, that is subject to significant externalities. Preservation of the resource provides worldwide benefits, but a local community may obtain value from depleting the resource. Conservation investors—typically represented by a non-governmental organization (NGO)—negotiate a quid-pro-quo contract by which the community foregoes destructive activities in exchange for payments from the NGO. Payments may be in the form of cash, services, or goods. Key aspects of these 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.

Conservation agreements are increasingly being adopted worldwide.[2] 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 (The Nature Conservancy 2013; Niesten and Gjertsen 2010; Svadlenak-Gomez et al. 2007). Experimentation with these programs is still fairly recent, as are preliminary assessments of their design and performance (Ezzine-de Blas et al. 2016; Honey-Roses et al. 2009; Niesten et al. 2008; Wunder et al. 2018; TNC and CI 2012).

There is a small theoretical literature on conservation agreements. Existing developments include analysis of moral hazard in conservation contract design (Ferraro 2008; Hart

[1] Miller and Watson (2013) and Watson (2013) developed theories of how contracting parties coordinate their behavior through bargaining and take advantage of inherent bargaining power. Prior work on transfers and renegotiation proofness includes: Goldlücke and Kranz (2012, 2013); Levin (2003); MacLeod and Malcomson (1989).

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

and Latacz-Lohmann 2005; Latacz-Lohmann and Van der Hamsvoort 1997; Wu and Babcock 1996), cost-effectiveness of conservation payments (Ferraro and Simpson 2002; Gibbons et al. 2011), marine conservation easements (Deacon and Parker 2009), and general conceptual models of payments for environmental services (Engel et al. 2008; Engel and Palmer 2008). However, none of these studies directly address the repeated nature of the interactions between the parties and the structure of self-enforcement (relational contracting).

We explore the incentive problems for conservation agreements by developing a relational-contracting model and then using the model to organize a discussion of three case studies. 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 some novel aspects in relation to the relational-contracting literature. Most prominently, informational asymmetry and incentive problems exist on both sides: The Community chooses whether to protect the resource and the NGO must engage in costly monitoring in order to obtain information about the Community's behavior. The model also allows for the parties to have different discount factors. Further, we examine extensions of the model that vary the monitoring technology and introduce the possibility of resource stock depletion.

Our model incorporates equilibrium selection and bargaining power by adopting 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. We characterize equilibrium play, including how punishments and rewards are structured, and we examine the relationship between the welfare attained in equilibrium and the parameters of the model, such as the parties' relative bargaining power, the monitoring technology, the benefits of preserving or exploiting the natural resource, and the discount factors. In the model, equilibrium joint value is increasing in the Community's bargaining power, which is also associated with a lower intensity of monitoring by the NGO.[3] We also describe how a 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.

We offer three case studies that differ in terms of success 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

[3]As first shown by Miller and Watson (2013), bargaining power plays a role because it is optimal for the parties to specify punishment paths that are less efficient than is their desired cooperation path, even though they anticipate renegotiating to achieve the joint value of cooperation. In equilibrium, the effective punishments depend on how the parties share in the surplus of renegotiation (dividing the surplus according to their bargaining powers), and incentives to cooperate depend on the severity of the punishments.

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 identified by the model and the broader theoretical literature. In particular, we highlight active contracting by the parties, instances of renegotiation, selection of imperfect monitoring, and the extent to which the agreements are documented.

By combining theory and case studies, we hope to connect repeated-game and contract theorists with scholars who study conservation agreements, emphasizing the theme of renegotiation. Conservation agreements are a particularly good class of relationships to examine with relational contract theory, for two main reasons. First, many conservation agreements have a bilateral structure, where it is clear who is involved in the negotiation process and what are the benefits of renegotiation. It may be easier to identify the effects of bargaining power in bilateral cases compared to settings with multilateral ongoing relationships, because the latter are complicated by the prospect of coalitions and side deals. By contrast, much of the prior theoretical literature on long-term environmental agreements and renegotiation looks at cases of multilateral relationships, such as Asheim and Holtsmark (2009) on renegotiation-proof equilibria in a multilateral emissions model. Second, conservation agreements by nature entail transfers of money or other resources as payments for environmental protection, so inherently these are settings with a degree of transferrable utility. Transfers simplify the equilibrium characterization in repeated-game models and they play an important role in our solution concept.[4]

Our case studies illustrate that renegotiation is an important force in real relationships, influencing both contractual arrangements and outcomes. They also suggest the real exercise of bargaining power, although the case studies provide insufficient data for differentiating between contractual equilibrium and other theories of renegotiation. More generally, the modeling exercise provides a useful backdrop and the concept of contractual equilibrium is particularly helpful because it has more predictive power relative to other concepts and is also more easily characterized.[5] The cases, in turn, show how real parties grapple with the elements essential for cooperation and sometimes are not successful, indicating limitations of the theory and perhaps pointing the way to future theoretical work.

[4]Overall, our modeling exercise contributes to the growing theoretical literature that utilizes repeated game models to examine self-enforced environmental agreements (examples include Asheim and Holtsmark 2009; Barrett 1994, 2005; Finus and Rundshagen 1998). 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 abstract notions of "renegotiation proofness" (e.g., Bernheim and Ray 1989; Farrell and Maskin 1989), in which bargaining power plays no role and negotiation is not modeled directly. Our use of contractual equilibrium adds structure and also simplifies the equilibrium characterization, which we hope will stimulate further work that is inspired by the complexities of real ongoing relationships.

[5]Most importantly, contractual equilibrium entails an explicit account of renegotiation and how it triggers a shift in the parties' coordinated expectations regarding future play—specifically, keeping track of what the parties would do, after any given history, in the event that they fail to renegotiate (even where the theory predicts that they will successfully renegotiate). We discuss this further in the next section.

In our model, the NGO's choice of whether to monitor can be regarded as a technology choice. Harstad, Lancia and Russo (2017) analyze a different type of technology choice within a period—one that affects the costs and benefits of emissions selected later in the period.[6] In another related vein, Harstad (2016) 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.[7]

The next section provides the details of our basic model and the solution concept. The equilibrium is characterized in Section 3, which also provides some simple comparative statics, the interpretation of equilibrium as a series of short-run contracts, and notes on how contracts may be documented. Section 4 presents two extensions of the model, one that allows the NGO within each period to commit to a monitoring level and one that addresses depletable resources. Section 5 contains the case studies. Section 6 concludes with some remarks about future work and applications. Appendix A contains some of the calculations behind the analysis in Section 3 and Section 4.

2 The Basic Model

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

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

Just before the bargaining phase, the parties observe the outcome of an arbitrary public randomization device, which allows them to randomize over how to coordinate their behavior in the current and future periods. Monetary transfers in the bargaining phase will typically go from the NGO to the Community, and so we let m be the amount transferred in this direction (a negative number if the net transfer goes in the other direction).

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

[6] Harstad, Lancia and Russo (2017) 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. Ramey and Watson (1997) examine how a long-term technology choice affects incentives to cooperate in a relational contract.

[7] In Harstad (2016), 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.

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

The Community chooses to either "protect" (P) or "exploit" (E) its local natural resource. If the Community exploits, then it obtains a gain of $e > 0$ in the period. If the Community protects, 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 $c > 0$.

Payoffs in a period are the sum of transfers and the stage-game payoff. Let $u = (u_C, u_N)$, where u_C and u_N denote, respectively, the Community's and NGO's stage-game payoff functions defined by the payoff table above. If transfer m is paid from the NGO to the Community in the negotiation phase, and action profile $a \in \{P, E\} \times \{M, R\}$ is played in the action phase, then the payoff vector in the current period is $(m, -m) + u(a)$.

The public signal represents a monitoring technology with fixed accuracy parameter $\lambda \in (0, 1]$ and possible signal realizations G (good) and B (bad). If the NGO selected M and the Community selected E in the period, then the signal is B with probability λ and G with probability $1 - \lambda$. Otherwise the signal is G for sure. We assume that the NGO's action and the monitoring signal are publicly observed, but the Community's action is private. Further, the NGO does not observe its own payoff. Therefore, only by monitoring can the NGO detect whether the Community chose P or E. The set of possible public outcomes in the action phase is $\Phi = \{MG, MB, RG\}$.

These informational assumptions 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. 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. The NGO may be able to estimate the Community's behavior in the long run by evaluating the health of the turtle population over time, but this assessment entails a large time lag and significant noise. If the NGO wants a signal of the Community's action within a period, then it will have to send a worker to the village to record an observation in person. The parameter λ accounts for possible error in the worker's performance. The model also captures cases in which the NGO hires a third party, such as an outside firm, to provide the monitoring service. In any case, our basic model assumes that the NGO's choice in a period is not observed by the Community until after the Community decides whether to protect or exploit.

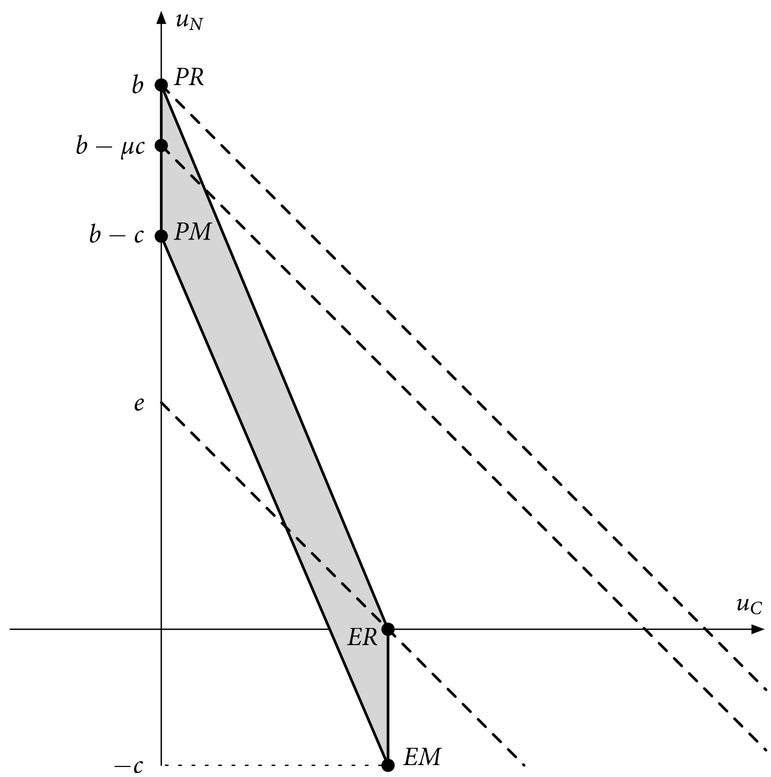

FIGURE 1. Stage game payoffs. The grey region is attainable without monetary transfers, using mixtures of action profiles PR, ER, PM, and EM. Heavy dashed lines illustrate payoff vectors that are attainable with monetary transfers combined with pure action profiles PR, ER, and a mixture of PM with probability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

Because the parties can make transfers, welfare is given by the sum of their payoffs, which we call their *joint value*. We assume that $b - c > e$, so that it is more efficient to have the action profile PM in the stage game, 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, because monitoring does not create a direct benefit to either party. Monitoring only generates information that randomly identifies an incident of exploitation. The structure of payoffs in this game is shown in Figure 1.

Our analysis of this repeated game will be stated in terms of *continuation values*. A party's continuation value from a given period is the discounted sum of its payoffs from this period. We allow for the parties to have different discount factors. Let δ_C and δ_N denote the discount factors of the Community and NGO, respectively; these are assumed to be strictly less than 1.

To consider incentives within a period, it is convenient to write each party's continuation value as the sum of its payoff within the current period and its discounted continuation value from the start of the next period. Letting v_C^t and v_N^t denote the Community's and NGO's continuation values from the start of period t, we call $v^t = (v_C^t, v_N^t)$ the "continuation value vector," or simply the "continuation value." If in a given period t the transfer is m^t and the action profile is a^t, then we have

$$v^t = (v_C^t, v_N^t) = \big(m^t + u_C(a^t) + \delta_C \mathbb{E}\big(v_C^{t+1}\big),\ -m^t + u_N(a^t) + \delta_N \mathbb{E}\big(v_N^{t+1}\big)\big).$$

Note that v_C^{t+1} and v_N^{t+1} are random variables from the perspective of period t because they can be conditioned on the random public outcome of interaction in period t. Incentives in period t depend on how this conditioning is structured. Importantly, given the absence of an effective legal system, there is no external enforcement, so any agreement must be self-enforced. That is, the parties can sustain cooperation only by appropriately rewarding or punishing each other over time.

2.1 Solution Concept

Our solution concept is a refinement of perfect public equilibrium called *contractual equilibrium* (Miller and Watson 2013; Watson 2013) where, in each period, (i) behavior in the action phase is consistent with individual incentives and (ii) the outcome of bargaining is given by the Nash bargaining solution with exogenous bargaining weights that represent a fixed bargaining protocol.[8] The disagreement point entails equilibrium play from the action phase of the current period and no immediate transfer. Importantly, a contractual equilibrium accounts for how disagreement play depends on the history of outcomes, and an agreement in one period implicitly specifies the disagreement points in future periods.

Bargaining must satisfy both *internal consistency* and *external consistency* conditions. The former means that the parties attain the maximum joint value that can be achieved from the current period, assuming that they utilize continuation values available within their current equilibrium. The latter means that they attain the maximum joint value over all internally consistent equilibria.

We define contractual equilibrium in terms of the set of continuation values attained in equilibrium, called the *contractual equilibrium value (CEV) set*, using a recursive formulation along the lines of Abreu, Pearce and Stacchetti (1990). At the end of this subsection we discuss how the characterization of the CEV set translates into a specification of behavior, and in the next subsection we offer an interpretation in terms of short-term contracts.

To simplify matters, we start by recognizing that the CEV set must be a compact line segment in $\mathbb{R}^2$ with slope -1. Formally this is established by Miller and Watson (2013)

[8] Miller and Watson (2013) provide noncooperative foundations for contractual equilibrium, where the cooperative bargaining solution described here is replaced with a noncooperative bargaining protocol and axioms constrain how cheap-talk messages are interpreted in equilibrium.

(Theorem 11 in the article's Supplement), but the intuition is simple: Because the parties can make transfers in the negotiation phase, they will always agree to obtain the maximum joint value that can be achieved; thus the equilibrium joint continuation value is constant across histories. Further, the parties can utilize the public randomization device to achieve continuation values between the extreme points.

So we know the CEV set has the form $V = \{\sigma z^C + (1-\sigma) z^N \mid \sigma \in [0,1]\}$, where the endpoints z^C and z^N satisfy

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

for some value L that we call the welfare level, or simply the *level*, of V. By convention we let z^C denote the endpoint that is worse for the Community and let z^N be the endpoint that is worse for the NGO. We define $d = z_C^N - z_C^C = z_N^C - z_N^N$ to be the *span* of V.

We next describe the conditions that characterize the CEV set. The first step is to describe payoffs that can be achieved *from the action phase* of a given period t assuming that the continuation value in period $t+1$ is an element of V. Because the parties can condition their future behavior on the public outcome in period t, their selection of a continuation value in period $t+1$ can be viewed as a function of the public outcome in period t. Letting $g\colon \Phi \to V$ denote this continuation value function, we write $g = (g_C, g_N)$. Let $\alpha = (\alpha_C, \alpha_N)$ denote a stage-game mixed action profile, where α_C is the Community's probability distribution over $\{P, E\}$ and α_N is the NGO's probability distribution over $\{M, R\}$. Let $\phi(\alpha)$ denote the resulting distribution over Φ.[9] Interaction in the action phase of period t is essentially a game with action profiles $\{P, E\} \times \{M, R\}$ and payoffs given by $u_C(\cdot) + \delta_C\, g_C(\phi(\cdot))$ for the Community and $u_N(\cdot) + \delta_N\, g_N(\phi(\cdot))$ for the NGO. We refer to this as the *g-induced game*.

Incentive compatibility requires mutual best-response behavior, i.e., a Nash equilibrium in the g-induced game. If the parties coordinate on action profile α in the action phase of period t, then the continuation value from the period-t action phase is

$$w = (w_C, w_N) = \big(u_C(\alpha) + \delta_C\, g_C(\phi(\alpha)),\ u_N(\alpha) + \delta_N\, g_N(\phi(\alpha))\big). \tag{2}$$

Therefore, the set of attainable continuation values from the action phase in period t is given by

$$D(V) \equiv \big\{ w \in \mathbb{R}^2 \;\big|\; \text{there is a function } g\colon \Phi \to V \text{ and a mixed action profile } \alpha \text{ that is a Nash equilibrium of the } g\text{-induced game, such that (2) holds} \big\}.$$

[9] For example, if α_C specifies P with probability $1/2$ and α_N specifies M for sure, then $\phi(\alpha)$ puts probability $\lambda/2$ on MB and probability $1 - (\lambda/2)$ on MG.

We next incorporate the bargaining theory to define the CEV set. Let π_C and π_N denote the exogenous bargaining weights for the Community and NGO, respectively, where $\pi_C, \pi_N \geq 0$ and $\pi_C + \pi_N = 1$. We write $\pi = (\pi_C, \pi_N)$.

Definition 1. A line segment $V = \{\sigma z^C + (1-\sigma)z^N \mid \sigma \in [0,1]\}$ with endpoints $\{z^C, z^N\}$ and level $L = z^C_C + z^C_N$ is a **bargaining self-generated (BSG) set** if

$$L = \max_{w \in D(V)} w_C + w_N \tag{3}$$

and, for each $j = C, N$, there exists a disagreement point $y^j \in D(V)$ such that endpoint z^j is the weighted Nash bargaining solution given level L and disagreement point y^j:

$$z^j = y^j + \pi \cdot (L - y^j_C - y^j_N).$$

A BSG set V is called a **contractual equilibrium value (CEV) set** if there is no other BSG set that attains a higher level.[10]

Note that each disagreement point y^j is a continuation value attainable from the action phase of the current period, which means it is a continuation value from the bargaining phase when no monetary transfer is made (as the solution concept requires). Internal consistency is represented by the BSG condition: In the negotiation phase, the players coordinate to maximize their joint continuation value from the current period, subject to the continuation values in the next period being consistent with the current equilibrium (these values are in V). The surplus of negotiation is this maximized joint value minus the joint value of the disagreement point. The parties divide the surplus according to their bargaining weights. External consistency is represented by selecting the BSG set with the highest level, meaning that the parties would renegotiate away from an equilibrium characterized by some BSG set with a lower level.

The CEV set has the longest span d among BSG sets, because longer spans enable higher powered incentives and thus help the parties attain higher payoffs. Accordingly, to construct the CEV set one should seek to maximize the span, subject to the BSG condition. This is how we proceed in the next section, where the CEV endpoints are determined by the simultaneous solutions to three interrelated optimization problems. The theorems of Miller and Watson (2013) guarantee existence, uniqueness, and compactness of the CEV set in a more general environment.

While the CEV set V is merely a set of payoff vectors, it can be translated into a specification of equilibrium strategies. For instance, suppose that following some history

[10] If there are multiple sets satisfying these criteria, which by definition must all attain the same level, then the definition in Miller and Watson (2013) selects the largest such set, which is unique.

the parties are meant to achieve a "promised" continuation value of $v \in V$. Let α^* and g^* support, in the definition of D, the value of w that maximizes $w_C + w_N$ (as in Equation 3). The parties attain v by making a transfer m in the bargaining phase, playing α^* in the action phase, and then, after realizing the public outcome $\phi \in \Phi$, continuing to the next period with a new promised continuation value of $g^*(\phi)$. Neither party can gain by deviating from α^*, since it is a Nash equilibrium of the g^*-induced game. The transfer m is calibrated to deliver exactly v:

$$v = (m, -m) + u(\alpha^*) + \big(\delta_C \, g_C(\phi(\alpha^*)),\ \delta_N \, g_N(\phi(\alpha^*))\big).$$

Of course, v must also be the outcome of the weighted Nash bargaining solution relative to an attainable disagreement value. We can always find a suitable disagreement value expressed as the expectation of a random draw over the disagreement values associated with the endpoints z^C and z^N. Therefore, it suffices to characterize behavior that achieves these extreme disagreement points and behavior that achieve the maximized joint value.

2.2 Comparison with other equilibrium concepts

For the analysis of recurrent negotiation, the leading alternatives to contractual equilibrium are the various definitions of *renegotiation-proofness*. These concepts define internal and external consistency in terms of "Pareto-perfection," which rules out Pareto-ranked continuation values.[11] Thus, renegotiation-proofness captures the idea that the parties, given the opportunity to renegotiate in each period, would not continue a specified path of play if there is another equilibrium path that they both prefer. Because these theories do not include an explicit model of the renegotiation process, they cannot distinguish between histories in which the parties failed to renegotiate and histories in which they renegotiated successfully. In particular, they don't contemplate what would happen if parties fail to reach an agreement in a period. Bargaining power and surplus therefore play no role. Contractual equilibrium, in contrast, is founded on a fuller account of the negotiation process, so that play under disagreement and play under agreement can be distinguished. Importantly, disagreement play may be less efficient than behavior under agreement.

It is not our objective to test contractual equilibrium against renegotiation-proofness or other solution concepts. Rather, the theory is meant to generate insights that help organize

[11] Bernheim and Ray (1989) and Farrell and Maskin (1989) give definitions of renegotiation-proofness (incorporating slightly different forms of Pareto-perfection) for standard repeated games; Goldlücke and Kranz (2012, 2013) provide definitions and analysis for settings with separate transfer and action stages in each period, as we have here. In some applications in the literature, Pareto-perfection is imposed at the transfer stage (see, for instance, Levin 2003); Fong and Surti (2009) impose Pareto-perfection at both stages. Renegotiation-proofness can be expressed in a recursive formulation, similar to how we characterize contractual equilibrium.

our thinking about, and evaluate the performance of, real conservation agreements. Some of the results we emphasize are consistent with contractual equilibrium and renegotiation-proofness, whereas other results are special to contractual equilibrium. We find contractual equilibrium to be particularly useful because it has the most structure, it explicitly accounts for the active contracting that we see in reality, and it yields the sharpest characterization of behavior and welfare.

3 Equilibrium Characterization

We now characterize the contractual equilibrium for our model of the Community-NGO relationship. Conditions derived in the next two subsections are based on the presumption that the contractual equilibrium entails protection of the natural resource in agreement, which requires that the span d is large enough to accommodate the conditions. The final steps of the analysis, in the third subsection below, include checking whether this is the case. We assume throughout that $\delta_C b \geq \delta_N e \geq \delta_C c$; these inequalities simplify a few steps of the analysis in Appendix A but are not critical to our results.

3.1 Structure of agreement play

The first step is to determine the structure of agreement play, including the stage-game action profile α and the selection of continuation values (a function $g\colon \Phi \to V$) that generates the highest joint value that can be achieved in equilibrium (maximizing $w_C + w_N$ among $w \in D(W)$). This is what, in every period, the parties renegotiate to achieve. Agreement play can be expressed as a function of only parameters and the span d, because the span provides the latitude for rewards and punishments via continuation values. The immediate transfer m^t, which divides the joint value in relation to the disagreement point, is analyzed in the following subsection.

Supposing that it is possible to give the Community the incentive to protect (P) rather than exploit (E) in the action phase, it is clear that the best agreement for the parties has the Community choosing P for sure in the current period t. It is also clear that the Community would have this incentive only if its continuation value from period $t+1$ were sensitive to its choice of P or E, and therefore the NGO must monitor with positive probability in period t. Since monitoring is costly, it is optimal for the NGO to monitor with a probability $\mu \in (0, 1]$ that is calculated below. The NGO's continuation value from period $t+1$ must depend on whether it monitors, and in such a fashion as to make the NGO indifferent between monitoring (M) and resting (R) in period t.

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.

Given that the Community selects P, the NGO's expected continuation value from the action phase in period t when it monitors in this period is $b - c + \delta_N(z_N^C - x)$. The NGO's expected value of resting is $b + \delta_N(z_N^C - d)$. These must be equal for the NGO to be indifferent, 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 action phase in period t when it selects P in this period is

$$\delta_C \left[\mu\big(z_C^C + x\big) + (1-\mu)\big(z_C^C + d\big)\right].$$

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 parties coordinate in period $t+1$. The Community's expected value of choosing E is

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

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 $\mu \geq e/\delta_C\lambda x$. 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 welfare level L in the agreement, it must be that L is the sum of the Community's and NGO's continuation values from the action phase, which is the same as the joint value from the bargaining phase (because the transfer nets out in the joint-value calculation). Adding the continuation values shown above yields:

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

3.2 Structure of disagreement play to support the CEV endpoints

Having characterized, as a function of d, the behavior under agreement and the welfare level, the next step is to characterize the disagreement points in $D(V)$ that correspond to the endpoints of the CEV set. That is, we know there are points $y^C \in D(V)$ and $y^N \in D(V)$, such that

$$\begin{aligned} z^C &= y^C + \pi\left(L - y_C^C - y_N^C\right), \\ z^N &= y^N + \pi\left(L - y_C^N - y_N^N\right). \end{aligned} \tag{7}$$

Along the way, we must characterize the stage-game behavior (α) and continuation-value selections (g) that yield values y^C and y^N.

The construction will nail down the equilibrium value of d. The key observation is that y^C and y^N are extreme points in the sense that, from these disagreement points, the parties negotiate to the endpoints of the CEV set. Because the outcome of negotiation is always in the direction of $\pi = (\pi_C, \pi_N)$ from any disagreement point, y^C must the be point in $D(V)$ that is furthest in the (perpendicular) direction $(-\pi_N, \pi_C)$ and y^N must be the point in $D(V)$ that is furthest in the opposite direction $(\pi_N, -\pi_C)$.

As shown in Appendix A, the disagreement point that most favors the Community, $y^N = (y_C^N, y_N^N)$, 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 the parties should disagree then the Community exploits the natural resource and the NGO does not monitor. This conclusion, which does not depend on d, implies:

$$y^N = \big(y_C^N, y_N^N\big) = \big(e + \delta_C z_C^N,\ \delta_N z_N^N\big). \tag{8}$$

Appendix A also shows that, for sufficiently large d, the disagreement point that most favors the NGO, $y^C = (y_C^C, y_N^C)$, involves playing PM in period t, followed by continuation value $z^C + (e/\lambda\delta_C,\ -e/\lambda\delta_C)$ from the start of period $t+1$.[12] 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

[12]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. A large enough span is required to provide room to punish the NGO if it does not monitor. This is addressed later.

probability 1. This specification yields:

$$y^C = (y_C^C, y_N^C) = \left(\delta_C z_C^C + \frac{e}{\lambda},\ b - c + \delta_N z_N^C - \frac{\delta_N e}{\delta_C \lambda}\right). \tag{9}$$

3.3 Completing the equilibrium characterization

To finish calculating the contractual equilibrium value set, we just need to combine equations and perform some algebra. Using equation 8 (which is two scalar equations) to substitute for y_C^N and y_N^N in the equation for z_C^N from (7), and also using Equation 1 to substitute for z_N^N, we obtain

$$z_C^N = \frac{\pi_C\big(1 - \delta_N\big)L + \pi_N e}{1 - \pi_N \delta_C - \pi_C \delta_N}. \tag{10}$$

Similarly, Using equation 9 to substitute for y_C^C and y_N^C in the equation for z_C^C from (7), and also using Equation 1 to substitute for z_N^C, we obtain

$$z_C^C = \frac{\pi_N \frac{e}{\lambda} + \pi_C \frac{\delta_N e}{\delta_C \lambda} + \pi_C\big(1 - \delta_N\big)L - \pi_C(b - c)}{1 - \pi_N \delta_C - \pi_C \delta_N}. \tag{11}$$

Let $\hat{d}$ denote the difference between z_C^N given by Equation 10 and z_C^C given by Equation 11. After simplifying terms and using $\pi_C + \pi_N = 1$, we have

$$\hat{d} = \frac{e - \pi_N \frac{e}{\lambda} - \pi_C \frac{\delta_N e}{\delta_C \lambda} + \pi_C(b - c - e)}{1 - \pi_N \delta_C - \pi_C \delta_N}.$$

Note that $\hat{d}$ is expressed in terms of only parameters.

Remember that our analysis is predicated on the presumption that the span of the CEV set is large enough to accommodate the conditions derived in the preceding subsections. Combining Equations 4 and 5 yields

$$d = \frac{c}{\delta_N} + \frac{e}{\delta_C \lambda \mu}.$$

Because μ is a probability and thus must be bounded above by 1, we obtain a lower bound on the span, written here as a condition on $\hat{d}$:

$$\hat{d} \geq \frac{c}{\delta_N} + \frac{e}{\delta_C \lambda}. \tag{12}$$

The preceding analysis proves the following result.

Theorem 1. *If Inequality 12 holds, the contractual equilibrium is characterized as follows. The span of the CEV set is given by $d = \hat{d}$ and the values of x, μ, L, and the CEV endpoints*

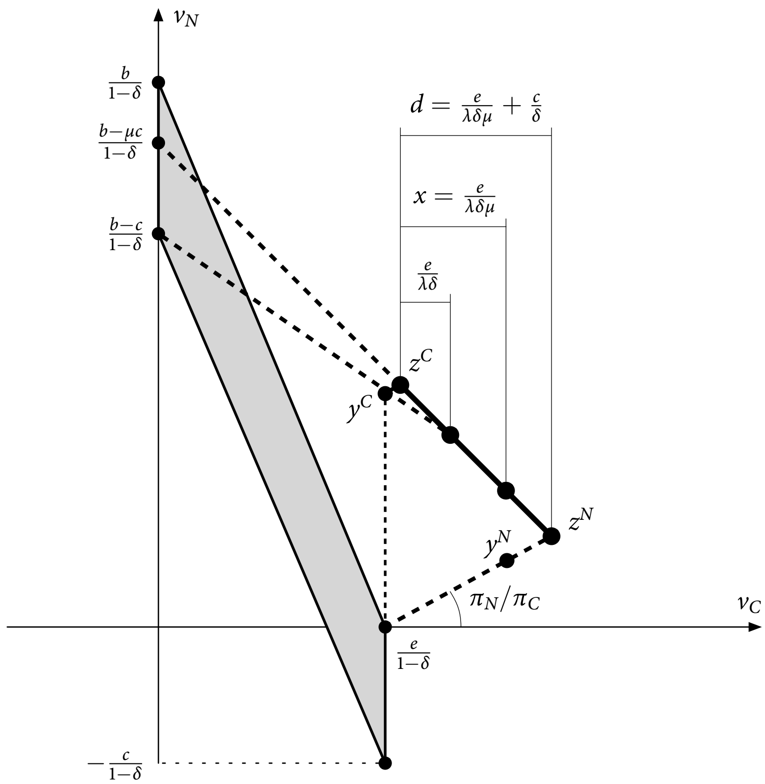

FIGURE 2. Contractual equilibrium for the case $\delta_N = \delta_C = \delta$. All payoffs are shown in total discounted terms. The contractual equilibrium value set V is attained along the equilibrium path. Its endpoints, z^C and z^N, are the expected payoffs in the states C and N, respectively. The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

z^C and z^N are calculated by using, in order, Equations 4, 5, 6, 10, 11, and 1. The parties are able to protect the natural resource in agreement, with random monitoring. Following a period t in which the parties are in agreement, the NGO's equilibrium payment to the Community in period $t+1$ depends on the public outcome of the stage game in period t and is increasing in the following order: RG, MG, MB.

If Inequality 12 fails then $d = 0$; the contractual equilibrium entails no protection of the natural resource and no monitoring (repeated play of the stage-game Nash equilibrium ER). In this case the CEV set is given by $z^C = z^N = (e/(1-\delta), 0)$.

The CEV set, extremal disagreement points, and equilibrium values x and μ are displayed graphically in Figure 2 for the case in which protection of the natural resource is possible and the parties have the same discount factor.

3.4 Documentation and interpretation as a series of short-run contracts

In principle the parties do not have to document their contractual arrangements, except to circumvent memory or communication barriers—that is, to keep track of how they are supposed to behave in various contingencies and to ensure that they have shared expectations.[13] In practice, parties routinely document their contracts. While our model makes no predictions about the extent parties will or should document their agreements, it is worth noting that contractual equilibrium can be interpreted as a series of short-term contracts that lends itself to a simple form of documentation.

In this form, the parties document (a) their intended behavior for the current period and (b) how their behavior in the next period, including an expectation to renew the agreement, should depend on the public outcome in the current period. Our model suggests two alternative ways in which the specification of behavior for the following period may be expressed. The first is to document the disagreement play that the parties would coordinate on in each contingency, with an understanding that renegotiation will occur in some contingencies (so the renegotiated transfer in the next period is implied). The second is to describe the monetary rewards and punishments that, in the model, would result from the renegotiation process. Implicit in this description is the disagreement play that the parties would resort to in the event that either would be unwilling to make the prescribed transfer at the beginning of the next period.

As an illustration of the second alternative, below is a list of what such a short-term contract would specify. In italics is a description of the disagreement play that is implicit in the agreement, which is what the parties would coordinate on in the event that the NGO refuses to pay the specified amount or the parties otherwise fail to renew the contract.

- The Community agrees to protect the natural resource, and the NGO agrees to monitor at a specific level (μ^* in the model).
- If the NGO does not monitor in the current period, then in the next period it must pay a large amount m^{RG} to the Community and the parties renew the contract. *Disagreement play: exploitation and no monitoring until the parties renegotiate to restore cooperation.*
- If the NGO monitors and obtains evidence that the Community exploited the resource, then the NGO must pay a small amount m^{MB} and the parties renew the contract. *Disagreement play: protection and monitoring for sure until the parties renegotiate to restore cooperation.*
- If the NGO monitors and there is no sign of exploitation, then the NGO must pay an intermediate amount m^{MG} and the parties renew the contract. *Disagreement play:*

[13] If the parties were able to take advantage of external enforcement, which is not the case here, then documentation may be needed to communicate their agreements to the third-party enforcer.

randomization between the disagreement play described in the previous two contingencies.

As described in Theorem 1, the specified payments satisfy $m^{RG} > m^{MG} > m^{MB}$. The Community implicitly shares in the cost of monitoring by receiving less when monitored (with signal G) than when not monitored, which follows from the NGO's required indifference condition.

3.5 Comparative statics

In this subsection, we provide results on how the contractual equilibrium values vary with the parameters of the model. These results are proved by taking derivatives of the expressions calculated above.[14] The first of these results summarizes what we can say in general, without restricting the parameters.

Theorem 2. *Assume that Condition 12 for a non-degenerate contractual equilibrium holds. Then the contractual-equilibrium span d is increasing in b and λ, and d is decreasing in e and c. If the parties are sufficiently patient and the monitoring is sufficiently precise (δ_N, δ_C, and λ are sufficiently high), then d is also increasing in both 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 12 for a non-degenerate contractual equilibrium holds. Then the contractual-equilibrium span is*

$$d = \frac{e\left(\frac{\lambda-1}{\lambda}\right) + \pi_C(b-c-e)}{1-\delta},$$

the equilibrium monitoring probability is $\mu = e/\lambda(\delta d - c)$, and the welfare level is $L = (b-\mu c)/(1-\delta)$. 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 they are decreasing in e and c.

The comparative statics with respect to π_C and $\pi_N = 1-\pi_C$ are illustrated in Figure 3. 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.

[14] Detailed calculations are available on request.

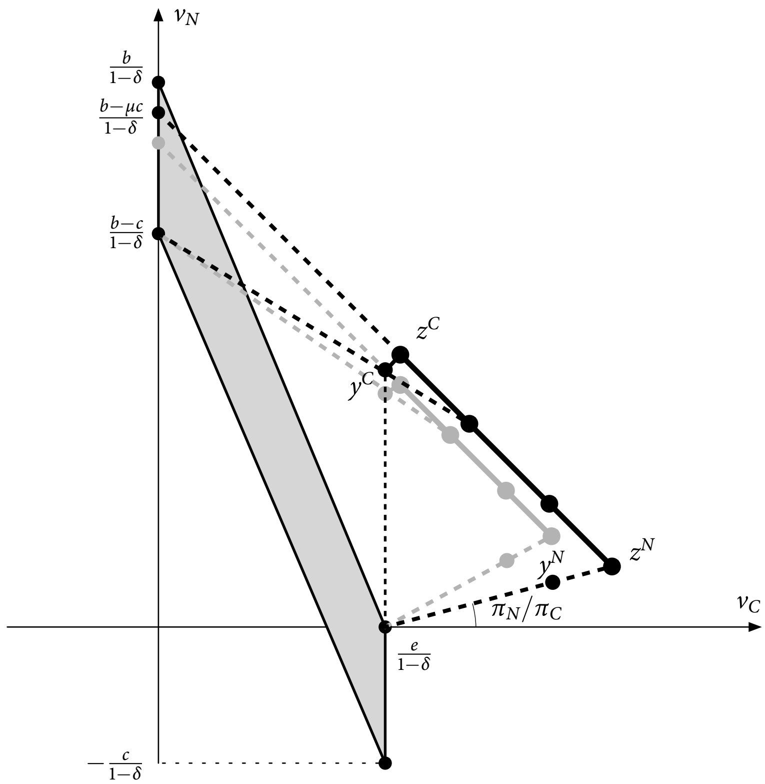

FIGURE 3. Comparative statics with respect to π_N/π_C, for the case $\delta_N = \delta_C = \delta$. As π_N/π_C decreases, the contractual equilibrium moves from the grey construction to the black construction. Because y^N is farther below the equilibrium joint value than is y^C, the increase in the Community's 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.

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 parties, changes in relative bargaining power have a greater influence on the endpoint most favoring the Community. Therefore the parties jointly 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.[15] 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.[16]

4 Variants of the Basic Model

This section describes two variants of the basic model, one dealing with the monitoring technology and another introducing a resource stock.

4.1 Commitment to third-party monitoring

In our basic model, the NGO's monitoring choice in a given period is not observed by the Community until after the Community chooses whether to exploit or protect. In some real cases, particularly where the NGO contracts with a third party to perform the monitoring function, the NGO can commit to a level of monitoring before the Community acts. To study how such commitment alters incentives in the relationship, in this subsection we examine a variant of the model in which the stage game is sequential: First, the NGO hires an agent to monitor at some frequency p (chosen by the NGO). The Community observes p and then chooses whether to exploit or protect. Finally, the agent selects M with probability p and R with probability $1 - p$, as contracted.[17] Let us call this the "commitment setting" and, to keep things simple, assume the parties have the same discount factor.

We assume that the monitoring service is competitively supplied. Therefore, because the act of monitoring entails cost c, the NGO must pay cp to the agent (as an up-front payment or in expected terms) to obtain monitoring frequency p. Payoffs and the signal distribution within a period are as in the basic model. If the Community selects E then its payoff is e and the NGO's payoff is $-cp$. If the Community selects P then its payoff is 0 and the NGO's payoff is $b - cp$. The NGO does not observe the Community's action or its own payoff. The monitoring signal is G for sure if the Community selects P and/or the

[15] To see this graphically, take $\pi_N/\pi_C \to \infty$ in Figure 3. As the line through $\big(e/(1-\delta), 0\big)$ 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.

[16] 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.

[17] We assume that the monitoring frequency can be enforced, which requires that either the agent's choice between M and R, or the agent's mixing probability, can be verified.

agent selects R. If the agent selects M and the Community selects E, then the signal is B with probability $\lambda > 0$ and G with probability $1 - \lambda$.

The contractual equilibrium in the commitment setting is similar to that of the basic model. However, cooperation is easier to sustain, and the contractual arrangements between the NGO and Community are simpler. The key observation is that, on the agreement path, the NGO's stage-game incentive constraint is relaxed compared to the corresponding constraint in the basic model. Before, parties needed to coordinate on continuation values to make the NGO indifferent between M and R. Because the NGO would save c by selecting R, the NGO's continuation value following MG needed to be exactly c/δ more than the NGO's continuation value following RG. The requirement in the commitment setting is just that NGO prefers selecting the contractual-equilibrium frequency p^* ex ante rather than deviating to $p = 0$ (the best deviation), which is weaker in two ways. First, if the NGO deviates then it saves only cp^* rather than c. Second, the Community begins the punishment by exploiting in the current period, so the NGO loses b immediately.

Here are some of the analytical details. In an agreement, conditional on the NGO choosing the prescribed frequency p^*, it is optimal for the parties to coordinate on continuation value z^C following MB, on $z^C + (x, -x)$ following MG, and on $z^C + (x', -x')$ following RG, where $x \in [0, d]$ and $x' \in [0, d]$ are to be determined. The Community's incentive condition is unchanged from the basic model, so Equation 5 applies. The NGO's incentive constraint (for the selection of p^*) is

$$-m^t + b - cp^* + \delta\big[p^*\big(z_N^C - x\big) + (1 - p^*)\big(z_N^C - x'\big)\big] \geq -m^t + \delta\big(z_N^C - d\big),$$

The left side is the NGO's continuation value from the beginning of period t when it chooses the prescribed monitoring frequency p^*. The right side is the value of deviating to $p = 0$, which results in play of E in the current period and continuation value $z^N = z^C + (d, -d)$ from period $t + 1$ (to punish the NGO). The number x' does not enter the Community's incentive condition and lowering it relaxes the NGO's condition, so it is best to set $x' = 0$. The NGO's incentive condition then simplifies to

$$x \leq \frac{d}{p^*} - \frac{c}{\delta} + \frac{b}{\delta p^*},$$

which is weaker than in the basic model where $x = d - (c/\delta)$. In fact, the right side exceeds d since $b > c$, and so the NGO's incentive condition is implied by the feasibility condition $x \leq d$. Because the level is decreasing in μ, it is optimal to set $x = d$ and then $p^* = e/\delta\lambda d$.

Appendix A shows that, as in the basic model, the disagreement point that most favors the Community involves playing $p = 0$ and E in period t, followed by continuation value z^N from the start of period $t+1$, regardless of the actual outcome in period t. The disagreement point that most favors the NGO involves playing exactly as in agreement, which means that renegotiation would not occur in this case. Thus, the contractual arrangement is simpler than in the basic model.

Theorem 4. *Assume that Condition 12 for a non-degenerate contractual equilibrium in the basic model holds. Relative to the basic model, contractual equilibrium in the commitment setting entails a strictly larger welfare level L, a strictly larger span d, and a strictly lower monitoring frequency ($p^* < \mu$). The contractual arrangement is simpler, in that, to punish the Community, disagreement play entails renewing the agreement rather than maximal monitoring. Also, following a period t in which the parties are in agreement, the NGO makes a payment to the Community in period $t+1$ only if monitoring reveals no sign of exploitation.*

4.2 Stock dynamics

For simplicity our basic 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. 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. 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 resource worthless to both parties, so their continuation values would then be zero.[18]

In this simple stock extension, 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, Equation 9 becomes

$$y^N = \left(e + \delta\beta z_C^N,\ \delta\beta z_N^N\right),$$

[18] In a more general setting of stock dynamics, the stock would be represented by a state variable θ, a transition process would be specified, and the set of contractual-equilibrium continuation values V would be a function of θ. Because the parties can negotiate and make transfers at the beginning of each period, $V(\theta)$ would not depend on the history except through θ and would be a line segment with slope -1.

because with probability $1-\beta$ the stock crashes, leading to continuation values of zero. Second, regarding agreement play, the Community's value of deviating to exploit the resource (that is, choosing E) becomes

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

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

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

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 parties. The latter effect is weak if π_C is relatively large. Thus, we obtain the following result. See Appendix A for details.

Theorem 5. *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.

5 Case Studies

In this section we present three cases of actual conservation agreements. 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. While no handful of case studies can be interpreted as statistical evidence in favor of any theory, these case studies offer suggestive support for our modeling approach, as they exhibit some of the phenomena predicted by contractual equilibrium.

It is worth noting first that the following technological conditions are required for a successful conservation agreement, in any equilibrium theory of behavior: First, the NGO has access to a monitoring technology that can detect whether the Community exploits the resource.[19] Second, the value of the resource to the NGO (representing world interests), net of the cost of monitoring the Community with great enough frequency, is greater than the Community's exploitation value. Third, 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. Finally, both the NGO and the Community are sufficiently patient. Additionally, in a contractual equilibrium, cooperation requires the Community's bargaining power to be high enough. Cases in the real world naturally vary in these parameters.

The modeling exercise indicates that, when evaluating real cases of conservation agreements, we should explore for signs of the following equilibrium elements in addition to the technical conditions:[20]

1. Active contracting between the Community and the NGO, specifying transfers in exchange for conservation services;
2. Recognition of the relational incentive problems;
3. Renegotiation and joint-value maximization;
4. Exercise of bargaining power;
5. Transfers that depend on whether exploitation was detected; and
6. History-dependent behavior in the absence of renegotiation.

On the first element, evidence that the parties have formed a contract could be documentation of the agreement or another demonstration of "meeting of the minds." The fifth element has several possible forms depending on the information technology and timing, as indicated by the modeling variations that we examined in Section 4. The sixth element means that disagreement play depends on which party should be rewarded and which punished given the recent history. Note that, since our model predicts agreement after any history, the sixth element describes a contingency that must be off the equilibrium path and therefore would not be observed if the parties were to always behave exactly as the solution concept prescribes. We proceed with the expectation that, for many reasons, behavior in reality will not always conform to an equilibrium strategy profile, yet the equilibrium theory still may provide guidance for interpreting behavior in what appears to be out-of-equilibrium contingencies.[21]

[19] 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.

[20] The second element is required by unrefined perfect-public equilibrium for any nontrivial equilibrium, elements 1, 3, and 6 are additional ones associated with renegotiation-proofness and contractual equilibrium, and elements 4 and 5 are unique to contractual equilibrium.

[21] Existing equilibrium concepts all specify what the parties will do out of equilibrium, following deviations, although deviations are, strictly speaking, inconsistent with the theory. In the real world, behavior rarely

Here is a synopsis of our case studies: All three cases feature active contracting and what appear to be earnest efforts to increase joint value. Renegotiation was observed in the first case following an apparent deviation, and clearly each party exercised significant bargaining power in restoring cooperation. The conservation agreements in the first and third cases were successful, while the second was unsuccessful. In all three cases the parties chose an imperfect monitoring regime, suggesting that they balanced the cost of monitoring against the incentive benefits (and realized that perfect monitoring is not needed). It appears that in these cases the monitoring technology affords a degree of commitment along the lines of our commitment modeling variant, because (i) we do not observe documentation of a direct link between payments and whether monitoring activity takes place, and (ii) in all three cases, third-party agents are involved in the monitoring process. Still, the communities share in the cost of monitoring to the degree that community members themselves engage in the monitoring activity; their monitoring effort can be regarded as in-kind payments to the NGOs. Finally, the first two conservation agreements were documented in ways that resemble the short-term arrangements that we have described, whereas long-term arrangements are documented more comprehensively in the third.

5.1 Forest protection in Cambodia

In Cambodia a conservation agreement was reached between the Commune Council of Chumnoab Commune (the "Community") and Conservation International Cambodia (the "NGO") to maintain and protect forest, wildlife, and crocodile habitat and to assist in combatting illegal hunting and wildlife trade.[22] The agreement was endorsed by the District Governor and the District Police chief. The initial agreement began in May 2006, with the understanding that the terms would be renewed on an annual basis indefinitely. The parties could also transition to a long-term agreement, which would involve reviewing terms periodically but not necessarily every year.

The NGO and Community developed a Participatory Land Use Plan (PLUP) that mapped where various activities were allowed or prohibited. The Community agreed to

conforms perfectly to an equilibrium theory. This is true for a variety of reasons, including inattention, lack of control, mistakes, and failure to coordinate perfectly on an equilibrium strategy profile (even when parties have actively contracted). To the extent that off-path considerations regarding deviations and beliefs are based on the idea of trembles or mistakes, equilibrium theory is applicable to off-path contingencies. Also, in line with the experimental literature on behavior in repeated games (Dal Bó and Fréchette Forthcoming; Dal Bó and Fréchette 2011), it may be appropriate to interpret the actors in our case studies as gradually learning to play as if in equilibrium, and it is the equilibria toward which they tend to converge that our theory is intended to describe. In this view, the parties should learn from the consequences of their "deviations" and "disagreements," as well as from observing others involved in similar relationships.

[22]The following sources provided background information for the case study discussed in this section: Conservation International (2007); Milne and Niesten (2009) and personal communication with Lykhim Ouk (Community Engagement Manager, CI-Cambodia).

follow the rules laid out in the plan, and to inform the Commune Natural Resource Management Council (CNRMC) of any observed or rumored hunting and wildlife trade activities. Community rangers agreed to patrol and report observed hunting and wildlife trade activities to the CNRMC. The CNRMC agreed to inform the NGO and the Forestry Administration of any violations.

In exchange for compliance with the agreement, the NGO would provide the community with benefits, including 8 water buffalo; funds for a teacher salary, construction of a school building, patrolling by the community rangers and police support; and funds for the CNRMC to organize community ranger patrols. Summing up, approximately $8,760 USD was to be spent annually to protect 6,555 hectares of forest.

The Forestry Administration, an independent party, was tasked with monitoring compliance through monthly joint patrols with community rangers. More intensive daily patrols of sensitive areas occurred during breeding seasons. A larger research team would conduct a comprehensive annual assessment. One of the main observable variables was whether forest was cleared.

The sanctions 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. Violations and sanctions in the Chumnoab Agreement

Transgressions	Sanctions
One or two families with water buffalo violate PLUP rules	Families lose water buffalo, and commune receives warning of 50% reduction of benefit package in the subsequent year.
Three or more families with water buffalo violate PLUP rules	Families lose water buffalo, and commune benefit package for the subsequent year reduced by 50%.
One or two families without water buffalo violate PLUP rules	Families go to bottom of list for receiving water buffalo, and commune receives warning of 50% reduction of benefit package in the subsequent year.
Three or more families without water buffalo violate PLUP rules	Families go to bottom of list for receiving water buffalo, and commune benefit package for the subsequent year reduced by 50%.

During the initial agreement period, the Community violated the contract by clearing forest. The Community initially claimed that the boundaries had not been clearly marked, but ultimately conceded that the agreement had been violated, warranting a sanction. The

Community and NGO agreed to a one-time waiver of the penalty, and to proceed with the following revised stipulations: (1) the Community would provide a list of names of the people responsible for the clearing and these families would be allowed to cultivate the cleared land for one season, after which the area would revert to protected status; (2) the NGO and the CNRMC would designate an additional, previously unprotected area elsewhere for protection; (3) the community would receive half the number of water buffalo.

A second agreement was entered into and the parties complied, which led to a new agreement the following year, which also achieved compliance. The level of monitoring effort fluctuated, particularly for the full annual assessment. In 2009, land reform affected the area, which led to changes in the project and rendered the agreement moot.

This case illustrates many of our model's key features. First, the basic conditions that the NGO be willing to make payments in excess of the value of the resources' exploitation value to the community and that the agreement be ongoing and renewable on an annual basis were met. In addition, they employed monitoring technologies that were able to detect the Community's deviations.

Beyond the technical conditions, this case exhibits signs of equilibrium elements consistent with our model. With reference to elements 1–6 on page 24, we summarize as follows.

1–2. The parties actively agreed on a contract that specifies payments in exchange for conservation effort, and describes how the Community will be punished in the event of noncompliance.

3–4. The model assumes that both parties exercise bargaining power, both initially and in every renegotiation. In this case, the parties settled the land clearing issue by agreeing to allow the Community to put additional forest area under protection in exchange for a reduction in punishment severity. The renegotiation thus involved both parties gaining compared to what the original contract specified, an indication that both have bargaining power. While the model does not predict violations on the equilibrium path, when interpreted as a series of short term agreements it predicts that the specified penalties will be renegotiated if violations are detected. Rather than follow 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.

5. The transfers to the Community that occurred after the first year of the agreement were reduced because exploitation was detected through monitoring, indicating that transfers are conditioned on monitoring outcomes. A key feature of our model is that since monitoring is costly, an optimal contract employs random monitoring, rather than constant monitoring. In this case an independent third party with constant

average monitoring costs is employed.[23] As shown in section 4.1, such a third party monitoring contract relaxes incentive constraints and can improve welfare.

6. Since every renegotiation in this case successfully reached agreement, we are unable to directly observe behavior in case of a failed renegotiation, as needed to verify feature 6 of the model. Nonetheless we surmise that if their renegotiation after the Community's deviation had failed then the penalties specified in the short term contract would have been implemented—that is, behavior under such a disagreement would depend on the prior behavior of the Community.

In summary, the Cambodia agreement illustrates a successful contract that follows the structure of our model and exhibits most of its key features. The next case will describe a conservation agreement that fails to account for some of the complexities of a repeated ongoing contractual relationship.

5.2 Laos deer conservation

The dry forest area in Savanakhet Province in Lao PDR is some of the last remaining habitat for the endangered Eld's deer. Hunting and habitat clearing by villagers living nearby was threatening the deer, despite legal protection by a wildlife conservation law since 1995. In 2003, the Wildlife Conservation Society and Smithsonian Institution (the "NGO") initiated a conservation payments program in Laos to increase the size of the Eld's deer population.[24] The Community agreed to maintain habitat by not expanding rice paddies, keeping cattle out of water holes, and establishing community patrols to report and stop poachers. In exchange, the NGOs agreed to an annual cash payment to each of three villages located near the deer habitat. The agreement stated that the NGOs would return at the end of each year to assess the deer population and make a payment equivalent to $300 USD contingent on an increase in the deer population. Villagers planned to use the payments for a village development fund and to pay per diems for meetings, patrolling, and education work by the Village Conservation Team.

At the end of the first year, the NGOs, government staff, and community conducted monitoring, but population estimates were questionable due to the small deer population. A habitat survey was included to assess whether the area used by deer had increased or decreased. According to the Wildlife Conservation Society, a shortage of funds and staff resources prohibited a more rigorous monitoring methodology. The monitoring suggested

[23] It also should be noted that this agreement includes some monitoring by community patrols that are paid for by the NGO. 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.

[24] 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 Stenhouse (2007).

that there was no change in the deer population, but the NGO decided to make the payment to the Community because villagers expressed excitement and support for the program.

In the second and third years the payment was increased to $450 USD to fund village development projects. 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 and is focused on land-use planning, villager-led patrolling of the sanctuary, and the development of sustainable livelihood opportunities.

With reference to the theory, this case meets some, but not all, of the technological conditions required for a successful conservation agreement. First, the NGO was willing to make payments in excess of the Community's value from poaching and habitat conversion. However, while the value of the resource to the NGO net of monitoring costs appeared to be greater than the Community's exploitation value, the lack of funding for adequate monitoring and community benefits may suggest otherwise. Second, both the NGOs and the Community clearly went into the initial agreement with an understanding that their arrangement would continue into future periods. Third, monitoring was planned to detect Community deviations. However, while the NGO had access to a monitoring technology to detect Community's exploitation, 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 to judge the compliance of the Community by whether this assessment showed a population increase. 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. Finally, the Community's bargaining power appeared to be high enough, given that they were able to receive payments despite exploitation.

Beyond the technical conditions, this case exhibits some signs of equilibrium elements consistent with our model but it seems to lacks others. With reference to elements 1–6 shown on page 24, we summarize as follows.

1–2. The parties actively agreed on a contract that specifies payments in exchange for conservation effort, and describes how the Community will be punished in the event of noncompliance. However, the agreement did not appropriately specify what should happen when the Community is caught exploiting and the parties fail to renegotiate.

3–4. While the model does not predict violations on the equilibrium path, when interpreted as a series of short term agreements it predicts that the specified penalties will be renegotiated if violations are detected. However, 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. The Community exercised its bargaining power, but the NGO did not appear to do so.

5–6. The most significant flaw of the implementation of this agreement is that the transfers did not depend on whether monitoring occurred and exploitation was detected. The NGO failed to anticipate that it would be unwilling to actually implement the punishment if the Community failed to comply with the agreement, and they were unable to renegotiate the agreement. Even when there was a clear violation of the agreement in its second year, the NGO was unwilling to withhold the payment as required by a strict enforcement of the agreement's provisions. The model suggests that in such an event the NGO should monitor much more closely while the Community protects the resource. 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.

This case highlights the challenge that many NGOs face in imposing conditionality on transfers to communities.[25] Once the organization has raised funds from donors and committed to protecting an area, it may face internal pressure to make payments despite infractions. The NGO may also fear retaliation from the Community that could cause long- term harm, even though the model predicts that the Community would choose to exploit the resource in such a situation and at the next opportunity negotiate with the NGO to restore cooperation. Practitioners in many conservation NGOs are uncomfortable with the concept of making payments for conservation in the first place. Withholding payments when conservation does not occur can be even less palatable, particularly when the funds are for a social good, such as education in poor villages. This issue of conditional aid has been controversial in the international development community for similar reasons (Doucouliagos and Paldam 2010; Paul 2015).

5.3 Grey whale habitat protection in Mexico

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 and hosts at least 221 other animal species, including numerous birds, green sea turtles, and bottlenose dolphins. In

[25] In a global dataset of 70 Payments for Environmental Services programs (of which our description of conservation agreements would be a subset), Wunder, Brouwer, Engel, Ezzine-de Blas, Muradian, Pascual and Pinto (2018) find that two-thirds monitored compliance comprehensively and the rest to some extent. Only 26 percent sanctioned non-compliance consistently, and 26 percent enforced the rules partially.

2005, the Laguna San Ignacio Conservation Alliance established a 120,000-acre conservation easement[26] comprising all the communal lands within the Ejido Luis Echeverria Alvarez on the southern shore of Laguna San Ignacio.[27] There are four parties to the agreement. The Ejido Luis Echeverria (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.[28] Maijanu is an organization that was created in the Community to receive and manage the funds disbursed through the easement.

The NGO conducts bi-annual monitoring of the area to determine compliance with the terms of the easement. An NGO team of biologists, GIS experts, and lawyers visit the same sites every six months, take photos, and compare them to originals. A trust fund disburses approximately $15,000 USD per year to conduct monitoring. The NGO also interviews 10-15 community members about whether they have noticed any changes. Community members also monitor throughout the year.

Each year, if the monitoring determines that the Community has met its obligations, the NGO (through ICF) agrees to pay approximately $25,000 USD to the Community (through Maijanu).[29] 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 then 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 payments will be withheld. If the violation created damage that can be restored, then the payments may be restarted once the damage is restored. If the damage cannot 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 the NGO can also take legal action to force compliance, which could include cessation of the illegal activity and restoration.

[26] A conservation easement is a 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.

[27] The following sources provided background information for the case study discussed in this section: Gjertsen and Niesten (2010), and personal communication with Raul Lopez (Ejido Luis Echeverria), Fernando Ochoa (Pronatura), Saul Alarcon (WildCoast), Ani Youatt (Natural Resources Defense Council), and Anne McEnany (International Community Foundation).

[28] ICF maintains a trust fund and manages it as a third party so there is transparency and accountability.

[29] This is the annual interest generated from the trust fund for community payments, which was capitalized in the amount of $650,000 USD. The NGO had planned for an increase in payments over time, but the Community chose to maintain a flat annual $25,000. As a result, the fund had grown to $808,000 USD as of 2017.

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

This case largely satisfies the technological requirements for cooperation. The grey whale habitat clearly has a high value to the public, given that the NGO was able to raise millions of dollars for its protection (an indication of the willingness to pay). 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. 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 being the case here, because monitoring costs are deducted from a separate account.

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 the trust fund for annual monitoring and payments. On bargaining power, 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.

Most of the equilibrium elements consistent with our model are present in this case. 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 irreversible damage were to occur, the parties should nonetheless find it optimal to renegotiate, in such a way that would punish the community while rewarding the NGO.[30]

While we regard the Laguna San Ignacio agreement as confirming the message of our modeling exercise regarding the ingredients essential for cooperation, it has a potentially important element that is outside our repeated-game model: a degree of external enforcement that may enhance incentives to cooperate, beyond what could be achieved by self-enforcement alone. In particular, some aspects of the contract may be enforceable in Mexican courts. Interaction between self-enforcement and external enforcement is an important topic for further study (see Watson, Miller and Olsen (2019) for recent work).

6 Conclusion

Our work contributes to the study of relational contracts by (i) applying the theory of contractual equilibrium to a principal-agent relationship with endogenous monitoring and renegotiation, and (ii) investigating how equilibrium behavior and the prospects of cooperation depend on the monitoring technology and other parameters. This paper demonstrates that the calculations required to characterize contractual equilibrium are straightforward, and we hope this will stimulate others to utilize the solution concept and consider variations.

Using the theory to organize our thinking about case studies puts us in the position of applying the theory to off-equilibrium contingencies. So when a community deviates from an agreement, as with forest protection in Cambodia, we ask whether the parties can renegotiate and strengthen their relationship as suggested by the theory, rather than interpreting a one-time deviation as a refutation of the theory. Similarly, when a conservation agreement fails, as with Laos deer conservation, we ask whether the apparent causes of the failure can be illuminated by our theory. Ultimately, while small numbers of case studies cannot be taken as statistical evidence in support of a theory, the case study approach provides details and documentation that can indicate whether the theory has the potential to positively explain and normatively guide. Future theoretical work may be usefully directed to model evolving sophistication in the context of relational contracts. That is, parties may "start small" in the sense of having a limited understanding of their strategic setting and incentives, adding levels of sophistication to their agreement as they encounter unanticipated events and work through problems (rather than deliberately starting small in stakes, as in Watson 1999, 2002 and Rauch and Watson 2003).

More modest theoretical steps would include examining general productive and monitoring technologies, outside options (in particular for the NGO), and resources with growth

[30]For instance, the Community could agree to allow the NGO to deduct a penalty amount from the trust fund, to spend on conservation efforts elsewhere.

and depletion dynamics. Furthermore, as in the case study 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 party 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.

Taking together our modeling exercise and case studies, we can offer a modest checklist to practitioners, which reinforces and expands the message from conservation scholars: (1) reaching an agreement requires that it generate enough surplus for the parties to share; (2) if it is possible to design the renegotiation process, then it should be done to endow the Community with sufficient bargaining power; (3) the NGO and Community should 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; (5) parties should determine how much monitoring is required to produce actionable information sufficient to provide incentives; and (6) the NGO should create the internal controls necessary to ensure conditionality of payments.

A Appendix

This appendix provides notes on the calculation of disagreement points and the extensions that were deferred from the main text. Much of the analysis here replicates the algorithm developed by Miller and Watson (2013) to characterizes the unique contractual equilibrium value (CEV) set V.

Disagreement play in the basic model

To prove that the stage game action profiles taken under disagreement for the extreme disagreement points are as described in Section 3, let us examine the incentive conditions and continuation values for a g-induced game, whereby $g(MG) = z^C + (\eta, -\eta)$, $g(MB) = z^C + (\eta', -\eta')$, and $g(RG) = z^C + (\eta'', -\eta'')$ for some numbers $\eta, \eta', \eta'' \in [0, d]$. Suppose that the parties coordinate on the mixed action profile in which the Community chooses P with probability α and the NGO chooses M with probability q, with the parties both indifferent between their two actions (which is required if $\alpha \in (0, 1)$ and $q \in (0, 1)$). The Community's indifference condition is

$$q\big[0 + \delta_C(z^C_C + \eta)\big] + (1-q)\big[0 + \delta_C(z^C_C + \eta'')\big] = \\ q\big[e + \delta_C(1-\lambda)(z^C_C + \eta) + \delta_C\lambda(z^C_C + \eta')\big] + (1-q)\big[e + \delta_C(z^C_C + \eta'')\big], \quad \text{(A1)}$$

where the left side is the Community's expected continuation value from the action phase when selecting P and the right side is the corresponding continuation value when selecting E. Solving for

η yields

$$\eta' = \eta - \frac{e}{q\lambda\delta_C}. \tag{A2}$$

The NGO's indifference condition is

$$\alpha\big[b + \delta_N(z^C_N - \eta'')\big] + (1-\alpha)\big[0 + \delta_N(z^C_N - \eta'')\big] = \\ \alpha\big[b - c + \delta_N(z^C_N - \eta)\big] + (1-\alpha)\big[-c + \delta_N(1-\lambda)(z^C_N - \eta) + \delta_N\lambda(z^C_N - \eta')\big], \tag{A3}$$

where the left side is the NGO's expected continuation value from the action phase when selecting R and the right side is the corresponding continuation value when selecting M. Solving for η'' and combining with Equation A2 yields

$$\eta'' = \eta + \frac{c}{\delta_N} - (1-\alpha)\frac{e}{q\delta_C}. \tag{A4}$$

This action profile and continuation-value selection implies the following continuation value from the action phase, calculated by using Equations A2 and A4 to subsitute for η' and η'' in the left sides of Equations A1 and A3:

$$(\tilde{y}_C, \tilde{y}_N) = \left(\delta_C z^C_C + \delta_C\eta + \frac{(1-q)\delta_C c}{\delta_N} - \frac{(1-q)(1-\alpha)e}{q},\ \alpha b + \delta_N z^C_N - c - \delta_N\eta + \frac{(1-\alpha)\delta_N e}{\delta_C q}\right).$$

Note that $\eta \in [e/q\lambda\delta_C, d]$.

Now we show that the extreme disagreement point in the direction $(\pi_N, -\pi_C)$, which is to punish the NGO and reward the Community, involves play of ER and selection of continuation value $z^N = z^C + (d, -d)$ regardless of the public outcome. This specification yields the following continuation value from the action phase:

$$(\overline{y}_C, \overline{y}_N) = \left(e + \delta_C z^C_C + \delta_C d,\ \delta_N z^C_N - \delta_N d\right).$$

A few algebraic steps show that $\overline{y}_C \geq \tilde{y}_C$ and $\overline{y}_N \leq \tilde{y}_N$ are implied by our assumptions that $b - c > e$ (in particular $b > c$) and $\delta_N e \geq \delta_C c$, and using the fact that $\eta \leq d$. This means that the disagreement point $(\overline{y}_C, \overline{y}_N)$ is furthest in the direction $(\pi_N, -\pi_C)$ compared to all possible equilibrium continuations in which the parties are indifferent between their two stage-game actions. It is easy to check that making one or both parties strictly prefer one of the stage-game actions can do no better.

We next show that the extreme disagreement point in the direction $(-\pi_N, \pi_C)$, which is to punish the Community and reward the NGO, involves play of MP and selection of continuation value $z^C + (e/\lambda\delta_C, -e/\lambda\delta_C)$ conditional on MG, z^C conditional on MB, and $z^N = z^C + (d, -d)$ conditional on RG. This specification yields the following continuation value from the action phase:

$$(\underline{y}_C, \underline{y}_N) = \left(\delta_C z^C_C + \frac{e}{\lambda},\ b + \delta_N z^C_N - \frac{\delta_N e}{\lambda\delta_C}\right).$$

Some tedious algebraic steps reveal that $\underline{y}_C \leq \tilde{y}_C$ and $\underline{y}_N \geq \tilde{y}_N$, where we use the assumption

$\delta_C b \geq \delta_N e$ in the second comparison.[31] This means that the disagreement point $(\underline{y}_C, \underline{y}_N)$ is furthest in the direction $(-\pi_N, \pi_C)$ compared to all possible equilibrium continuations in which the parties are indifferent between their two stage-game actions. As before, it is easy to check that making one or both parties strictly prefer one of the stage-game actions can do no better.

Disagreement play in the commitment setting

We follow the same lines to establish that the extreme disagreement points in the commitment setting are as described in Section 4, but the analysis is simpler because of the sequential moves in the stage game. Consider first the play from the action phase that achieves the disagreement point furthest in the direction $(\pi_N, -\pi_C)$, which is to punish the NGO and reward the Community. For any p, it is clearly best to have the Community play E and then proceed to continuation value z^N in the following period; play of P with the required continuation-value selection only raises the NGO's disagreement value and lowers the Community's. Stepping back to the choice of p, to achieve any strictly positive value requires rewarding the NGO in the next period, so the NGO's disagreement value would not decrease and the Community's disagreement value would decrease. Hence $p = 0$ is optimal.

Next consider the play from the action phase that achieves the disagreement point furthest in the direction $(-\pi_N, \pi_C)$, which is to punish the Community and reward the NGO. Fixing p, suppose the Community is to choose P with probability α. Considering the Community's incentive constraint, the best continuation-value selection that gives the Community the incentive to choose P with positive probability is to coordinate on $g(MG) = z^C + (e/\delta p\lambda,\ -e/\delta p\lambda)$ and $g(MB) = g(RG) = z^C$. Calculating the parties' continuation values from the action phase, the disagreement point is then

$$(y_C, y_N) = \left(\delta z_C^C + \frac{e}{\lambda},\ \alpha b - pc + \delta z_N^C - \frac{e}{\lambda} + e - \alpha e\right).$$

Clearly, increasing α causes y_N to increase and leaves y_C unchanged, so we want $\alpha = 1$ (the Community selects P for sure). Further, increasing p causes y_N to decrease and leaves y_C unchanged, so we want p to be as low as possible. Because $e/\delta p\lambda$ cannot exceed d (otherwise $g(MG)$ would not be in V), the optimal p is the minimal feasible number, $e/\delta d\lambda$, which is of course p^*.

Calculations for the simple stock extension

Here are some details for the proof of Theorem 5. Some algebraic manipulation reveals that

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

[31] Detailed calculations are available upon request.

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 $L^* = (b - \mu^* c)/(1-\delta)$, we also obtain that L^* rises.

References

Abreu, Dilip, David Pearce and Ennio Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58: 1041–1063.

Asheim, Geir B. and Bjart Holtsmark. 2009. "Renegotiation-Proof Climate Agreements with Full Participation: Conditions for Pareto-Efficiency." *Environmental and Resource Economics* 43: 519–533.

Barrett, Scott. 1994. "Self-Enforcing International Environmental Agreements." *Oxford Economic Papers* : 878–894.

Barrett, Scott. 2005. *The Theory of International Environmental Agreements*, volume 3 of *Handbooks in Economics*, chapter 28. Elsevier, pp. 1457 – 1516.

Bernheim, B. Douglas and Debraj Ray. 1989. "Collective Dynamic Consistency in Repeated Games." *Games and Economic Behavior* 1: 295–326.

Conservation International, Conservation Stewards Program. 2007. "Conservation Agreements: Field Guide for Design and Implementation." Arlington, Virginia.

Dal Bó, Pedro and Guillaume Fréchette. Forthcoming. "On the Determinants of Cooperation in Infinitely Repeated Games: A Survey." *Journal of Economic Literature* .

Dal Bó, Pedro and Guillaume R. Fréchette. 2011. "The Evolution of Cooperation in Infinitely Repeated Games: Experimental Evidence." *The American Economic Review* 101: 411–429.

Deacon, Robert T. and Dominic P. Parker. 2009. "Encumbering Harvest Rights to Protect Marine Environments: A Model of Marine Conservation Easements." *Australian Journal of Agricultural and Resource Economics* 53: 37–58.

Doucouliagos, Hristos and Martin Paldam. 2010. "Conditional Aid Effectiveness: A Meta-Study." *Journal of International Development: The Journal of the Development Studies Association* 22: 391–410.

Engel, Stefanie, Stefano Pagiola and Sven Wunder. 2008. "Designing Payments for Environmental

Services in Theory and Practice: An Overview of the Issues." *Ecological Economics* 65: 663–674.

Engel, Stefanie and Charles Palmer. 2008. "Payments for Environmental Services as an Alternative to Logging Under Weak Property Rights: The Case of Indonesia." *Ecological Economics* 65: 799–809.

Ezzine-de Blas, Driss, Sven Wunder, Manuel Ruiz-Pérez and Rocio del Pilar Moreno-Sanchez. 2016. "Global Patterns in the Implementation of Payments for Environmental Services." *PloS one* 11: e0149847.

Farrell, Joseph and Eric S. Maskin. 1989. "Renegotiation in Repeated Games." *Games and Economic Behavior* 1: 327–360.

Ferraro, Paul J. 2001. "Global Habitat Protection: Limitations of Development Interventions and a Role for Conservation Performance Payments." *Conservation Biology* 15: 990–1000.

Ferraro, Paul J. 2008. "Asymmetric Information and Contract Design for Payments for Environmental Services." *Ecological Economics* 65: 810–821.

Ferraro, Paul J. and Agnes Kiss. 2002. "Direct Payments to Conserve Biodiversity." *Science* 298: 1718–1719.

Ferraro, Paul J. and R. David Simpson. 2002. "The Cost-Effectiveness of Conservation Payments." *Land Economics* 78: 339–353.

Finus, Michael and Bianca Rundshagen. 1998. "Renegotiation–Proof Equilibria in a Global Emission Game When Players Are Impatient." *Environmental and Resource Economics* 12: 275–306.

Fong, Yuk-fai and Jay Surti. 2009. "The Optimal Degree of Cooperation in the Repeated Prisoners' Dilemma with Side Payments." *Games and Economic Behavior* 67: 277–291.

Gibbons, James M, Emily Nicholson, EJ Milner-Gulland and Julia PG Jones. 2011. "Should Payments for Biodiversity Conservation be Based on Action or Results?" *Journal of Applied Ecology* 48: 1218–1226.

Gjertsen, Heidi and Eduard Niesten. 2010. "Incentive-Based Approaches in Marine Conservation: Applications for Sea Turtles." *Conservation and Society* 8: 5–14.

Goldlücke, Susanne and Sebastian Kranz. 2012. "Infinitely Repeated Games with Public Monitoring and Monetary Transfers." *Journal of Economic Theory* 147: 1191–1221.

Goldlücke, Susanne and Sebastian Kranz. 2013. "Renegotiation-Proof Relational Contracts." *Games and Economic Behavior* 80: 157–178.

Hadri, K. and W. Mikhail, editors. 2014. *Econometric Methods and Their Applications in Finance, Macro and Related Fields (proceedings of the 15th Annual Conference of the African Econometric Society)*. World Scientific Publishing: Singapore.

Harstad, Bård. 2016. "The Market for Conservation and Other Hostages." *Journal of Economic Theory* 166: 124–151.

Harstad, Bård, Francesco Lancia and Alessia Russo. 2017. "Compliance Technology and Self-Enforcing Agreements." Working paper.

Hart, Rob and Uwe Latacz-Lohmann. 2005. "Combating Moral Hazard in Agri-Environmental Schemes: A Multiple-Agent Approach." *European Review of Agricultural Economics* 32: 75–91.

Honey-Roses, Jordi, Jose Lopez-Garcia, Eduardo Rendon-Salinas, Armando Peralta-Higuera and Carlos Galindo-Leal. 2009. "To Pay or Not to Pay? Monitoring Performance and Enforcing Conditionality when Paying for Forest Conservation in Mexico." *Environmental Conservation* 36: 120–128.

Latacz-Lohmann, Uwe and Carel Van der Hamsvoort. 1997. "Auctioning Conservation Contracts: A Theoretical Analysis and an Application." *American Journal of Agricultural Economics* 79: 407–418.

Levin, Jonathan. 2003. "Relational Incentive Contracts." *American Economic Review* 93: 835–857.

MacLeod, W. Bentley and James M. Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57: 447–80.

McShea, William. 2015. *Personal communication on October 20.* Smithsonian Institution.

Miller, David A and Joel Watson. 2013. "A Theory of Disagreement in Repeated Games with Bargaining." *Econometrica* 81: 2303–2350.

Milne, Sarah and Eduard Niesten. 2009. "Direct Payments for Biodiversity Conservation in Developing Countries: Practical Insights for Design and Implementation." *Oryx* 43: 530–541.

Niesten, E. and Heidi Gjertsen. 2010. "Economic Incentives for Marine Conservation." Technical report, Science and Knowledge Division, Conservation International, Arlington, Virginia.

Niesten, Eduard, Aaron Bruner, Rrichard Rice and Patricia Zurita. 2008. "Conservation Incentive Agreements: An Introduction and Lessons Learned to Date." Technical report, Conservation International, Washington, D.C.

Paul, Elisabeth. 2015. "Performance-Based Aid: Why It Will Probably Not Meet Its Promises." *Development Policy Review* 33: 313–323.

Ramey, Garey and Joel Watson. 1997. "Contractual Fragility, Job Destruction, and Business Cycles." *The Quarterly Journal of Economics* 112: 873–911.

Rauch, James E and Joel Watson. 2003. "Starting Small in an Unfamiliar Environment." *International Journal of Industrial Organization* 21: 1021–1042.

Simpson, R. David and Roger A. Sedjo. 1996. "Paying for the Conservation of Endangered Ecosystems: A Comparison of Direct and Indirect Approaches." *Environment and Development Economics* 1: 241–257.

Svadlenak-Gomez, Karin, Tom Clements, Charles Foley, Nikolai Kazakov, Dale Lewis, Dale Miguelle and Renae Stenhouse. 2007. "Paying for Results: WCS Experience with Direct Incentives for Conservation." Technical report, Wildlife Conservation Society, Bronx, NY.

The Nature Conservancy. 2013. "Fishery Buy-Ins: Applying Rights-Based Incentive Agreements to Sustainable Fishery Interventions." The Nature Conservancy, Narragansett, Rhode Island.

The Nature Conservancy and Conservation International. 2012. *Practitioner's Field Guide for Marine Conservation Agreements: Best Practices for Integrating Rights-based Incentive Agreements into Ocean and Coastal Conservation Efforts.* The Nature Conservancy, Narragansett, Rhode Island, final v2 edition.

Troëng, Sebastian and Carlos Drews. 2004. "Money Talks: Economic Aspects of Marine Turtle Use and Conservation." Technical report, WWF-International, Gland, Switzerland.

Watson, Joel. 1999. "Starting Small and Renegotiation." *Journal of Economic Theory* 85: 52–90.

Watson, Joel. 2002. "Starting Small and Commitment." *Games and Economic Behavior* 38: 176–199.

Watson, Joel. 2013. "Contract and Game Theory: Basic Concepts for Settings with Finite Horizons." *Games* 4: 457–496.

Watson, Joel, David A. Miller and Trond Olsen. 2019. "Relational Contracting, Negotiation, and

External Enforcement." Working paper.

Wu, JunJie and Bruce A Babcock. 1996. "Contract Design for the Purchase of Environmental Goods from Agriculture." *American Journal of Agricultural Economics* 78: 935–945.

Wunder, Sven. 2004. "The Efficiency of Payments for Environmental Services in Tropical Conservation." *Conservation Biology* 21: 48–58.

Wunder, Sven. 2008. "Payments for Environmental Services and the Poor: Concepts and Preliminary Evidence." *Environment and Development Economics* 13: 279–297.

Wunder, Sven, R Brouwer, S Engel, Driss Ezzine-de Blas, R Muradian, Unai Pascual and R Pinto. 2018. "From Principles to Practice in Paying for Nature's Services." *Nature Sustainability* 1: 145.