Coalitions and Recurrent Negotiation in Multilateral Relational Contracts (Preliminary)

David Miller and Joel Watson *

December 31, 2025

Abstract

We introduce *coalition-directed contractual equilibrium* to model equilibria in repeated games that are negotiated (and frequently renegotiated) by a coalition comprising a subset of the players (the *insiders*); the insiders' agreements also coordinate the equilibrium behavior of *outsiders* using promised rewards and punishments. However, because outsiders do not participate in negotiations and thus cannot claim a share of the surplus, insiders' promises to outsiders cannot extend beyond the next renegotiation. If renegotiation occurs each period, then such promises must take the form of direct payments from insiders to outsiders. In the coalition directed equilibrium of a repeated public-goods game, the coalition maximizes surplus by outsourcing marginal contributions to outsiders—that is, if an insider is otherwise identical to an outsider, then in equilibrium the outsider makes greater contributions, while the insider obtains a higher payoff. We interpret our results in light of global climate change negotiations.

*Thanks to Jim Malcomson, Paul Klemperer, Felix Vardy, Jack Fanning, and Lisa Bernstein. Additional thanks to seminar participants at Oxford, the 11th Workshop on Relational Contracts, and Michigan/Michigan State Theory Day.

1 Introduction

In settings in which multiple economic actors interact over time, a subset of these actors (a *coalition*) may convene to establish an agreement on how to behave (a *relational contract*). When no external authority exists to enforce their agreement, the coalition must rely on self-enforcement. Examples abound in the international sphere, including trade agreements, security agreements, peace treaties, and environmental agreements. A key economic consideration is the extent to which the coalition can influence the behavior the *outsiders* who are not members of the coalition.

We develop a repeated-game model to study how a coalition negotiates and renegotiates a relational contract in a setting without external enforcement. The coalition members care about, and seek to influence, the behavior of all players in the game, including the outsiders. Our key conceptual contribution is a new theory of coalitional power: that an agreement by the coalition coordinates *all* players on an equilibrium in the continuation of the game, subject to incentive constraints. In our model, the coordination role is the sole source of the coalition's power, as the productive technology does not depend on the composition of the coalition. However, the coalition's coordination power is limited, as it cannot commit not to renegotiate with some probability each period.

International environmental agreements (IEAs) are an important case of coalitional power. Most of the thousands of agreements formed in recent decades were made by a strict subset of countries acting as a coalition. Because these agreements deal with productive activity that entails externalities, such as global commons, the coalitions seek not only to regulate the behavior of coalition members but also the behavior of outsiders. The literature on IEAs has emphasized the idea that, by using trade sanctions and other methods to put outsiders at a disadvantage, the coalition may provide incentives for outsiders to join the coalition and thereby regulate their productive actions as signatories to the agreement.[1] Less

[1] For example, even as negotiations began to break down at the Conference of Parties (COP 15) to the United Nations Framework Convention on Climate Change in Copenhagen in 2015, the United States, Brazil, China, India, and South Africa negotiated the Copenhagen Accord in secret and then convinced most other nations to make greenhouse-gas abatement pledges under it even though it was not officially adopted by COP 15.

emphasized is a second possibility: inducing outsiders to behave more cooperatively without joining the coalition, such as through payments for environmental services or carbon offsets in service of the coalition's abatement program. More generally, some relational contracts are focused explicitly on influencing the behavior of outsiders who are excluded from the coalition; examples include security agreements such as the North Atlantic Treaty.

We show that the coalition's lack of commitment implies that it cannot motivate outsiders using rewards and punishments that extend beyond the next renegotiation. This holds for any subgame perfect equilibrium satisfying "coalitional internal consistency", a minimal collective-rationality condition that underlies most theories of collective rationality in games. First, suppose the coalition renegotiates in every period; then outsiders can be motivated only to best-respond within the period, as if they were short-run players. Our model allows for monetary transfers both before and after the productive stage-game actions. The coalition can promise payments to outsiders at the end of the period in return for taking cooperative actions. Coalition members use continuation rewards and punishments to motivate each other to make the required payments. Moreover, the coalition may be able to extort surplus from the outsiders by demanding transfers before the productive stage-game actions, by threatening to punish outsiders who fail to pay. We propose two different versions of coalitional internal consistency: one that allows for such extortion, and one that rules out extortion. If extortion is allowed, our result on how outsiders can be motivated holds regardless of the nature of the stage game, while if extortion is ruled out then the result holds if the stage game satisfies a concavity condition.

A central question is whether the coalition can induce outsiders to behave non-myopically, such as by exerting higher effort to abate their greenhouse-gas emissions than myopic free-riders would choose to. If the coalition must renegotiate in every period, then (under some relevant conditions on the stage game) the answer is negative: the coalition directs outsiders to simply best-respond in the stage game, since at the next period's renegotiation the coalition would renege on any promised punishment. If the coalition can make payments to outsiders at the end of the stage game, before the next renegotiation, then outsiders can be motivated to do more. If renegotiation occurs in each period with probability less than 1, then the coalition can promise dynamic rewards and punishments to further motivate outsiders, more strongly as the probability of renegotiation falls.

In order to obtain sharper results, we introduce "coalition-directed contractual equilibrium" (CDCE), which extends the concept of contractual equilibrium Miller and Watson (2013b) to accommodate coalitional power. In a CDCE, transfers made before the productive stage-game actions are assumed to be part of the coalition's negotiation process, so outsiders are not involved and cannot be extorted. We show that a CDCE is characterized by the fixed point of an aggregate incentive operator, which is straightforward to solve because its domain is one-dimensional. The fixed point of the operator is the aggregate incentive power of the equilibrium, which the coalition allocates optimally between incentivizing productive actions by insiders and incentivizing productive actions by outsiders.

We model a linear-quadratic public goods game to clarify the strategic forces at work when a coalition of nations negotiates an agreement on reducing greenhouse-gas emissions. We show that outsiders are directed to exert higher abatement effort than insiders (motivated by the promise of direct payments for emission reductions), and that insiders attain higher payoffs than outsiders.

2 The model

Our model is a repeated game with infinitely many periods $t \in \mathbb{N}$ and a finite set of players N, with $n \equiv |N|$. A coalition of players $C \subset N$ is fixed throughout the game. Players in the coalition are *insiders*, and those who are not are *outsiders*. The coalition is payoff-irrelevant; its influence derives only from its members' ability to reach agreements that coordinate the continuation play of both insiders and outsiders.

In each period, the players interact over several phases, as described below, including a phase in which the players choose productive actions. Monetary transfers are feasible, and payoffs are linear in these transfers. [2]

In each period, the players interact in the following phases in sequence:

[2] A public randomization device, normally assumed in repeated-game theory, would not add anything here due to the inclusion of voluntary end-of-period transfers, and so we leave it out of the model.

- *Negotiation phase:* With probability ρ, an opportunity arises for the coalition to negotiate, and the coalition then engages in cheap-talk bargaining over how to coordinate behavior in the continuation of the game. With probability $1-\rho$ there is no negotiation opportunity. Let $\phi \in \{0,1\}$ record whether an opportunity to negotiate occurred ($\phi = 1$) or not ($\phi = 0$). The negotiation phase is payoff-irrelevant.

- *Transfer phase:* Players simultaneously make voluntary transfers. The transfer of each player $i \in N$ is a vector

$$m^i \in \mathbb{T}^n(i) \equiv \left\{ \tau \in \mathbb{R}^n \;\middle|\; \tau_i \leq 0 \text{ and } \sum_{j \in N} \tau_j \leq 0 \right\}.$$

  where m^i_i is the total amount that player i gives up, m^i_j is the part given to player j for each $j \neq i$, and the remainder is thrown away. The resulting total transfer vector is $m = \sum_{i \in N} m^i \in \mathbb{R}^n$, and $m_i \in \mathbb{R}$ is the net transfer that accrues to player i.

- *Action phase:* Players simultaneously choose productive actions, with each player $i \in N$ choosing an action a_i from a finite set A_i. The set of action profiles is denoted by $A = A_1 \times A_2 \times \cdots A_n$. Players' receive payoffs given by function $u \colon A \to \mathbb{R}^n$, where $u_i(a)$ is the payoff of player i.

- *Payment phase:* Players simultaneously make voluntary payments. The payment of each player $i \in N$ is a vector $p^i \in \mathbb{T}^n(i)$, resulting in a total payment vector of $p = \sum_{i \in N} p^i$, and p_i is the net payment that accrues to player i.

Everything is publicly observed. The payoff vector for the period is $m + u(a) + p$. The players discount the future according to a common discount factor $\delta \in [0,1)$. Writing m^t, a^t, and p^t as the total transfer, action profile, and total payment in period t, respectively, (with apologies for dual use of superscripts on the transfer and payment), the discounted continuation payoff from the start of any period T is $\sum_{t=T}^{\infty} \delta^{t-T} (1-\delta)(m^t + u(a^t) + p^t)$. Setting $T = 1$ gives the payoff for the entire game.

It is useful to keep in mind two substantive applications to which we will return later in the paper. In the first application, the players are countries, and the coalition is the subset that makes an agreement on climate-change abatement. In the action phase, $a_i \in A_i = \mathbb{R}_+$ is country i's abatement effort, and $u_i(a)$ is the benefit country i receives from all countries' efforts net of the cost it incurs for its own effort. We assume u_i is increasing in a_j for all $j \neq i$, and that for all a_{-i} country i's production-phase best response is finite. In the second application, the players are the individuals in a business partnership, and the coalition is the set of managing partners. In the action phase, $a_i \in A_i = \mathbb{R}_+$ is individual i's productive effort, and $u_i(a)$ is individual i's share of revenue (a function of everyone's effort) minus her personal cost of her own effort, in monetary terms.

Base game and equilibrium concept

The foregoing fully describes a noncooperative game, except that we have not specified the bargaining protocol in the negotiation phase. But regardless of the bargaining protocol, since bargaining is an exercise in cheap talk with no direct payoff implications, there will generally be "babbling" subgame perfect equilibria in which the players ignore their past actions in the negotiation phase. Limiting attention to these equilibria is equivalent to studying a repeated game without active negotiation, where the negotiation phase consists of only the publicly observed random draw ϕ. Let us call the repeated game without active negotiation the *base game*. Throughout this paper we limit attention to pure-strategy equilibria.

There are two main approaches to formally introducing active negotiation into the model. The first approach is fully noncooperative: specify a noncooperative bargaining protocol describing how the players communicate in the negotiation phase, and impose axioms that relate their communication to their coordinated behavior in the continuation of the game. The axioms select among the subgame-perfect equilibria of this fully noncooperative game. The second approach is hybrid: assume that negotiation is resolved according to a cooperative theory of bargaining (axiomatic or otherwise), which effectively selects among the subgame-perfect equilibria of the base game.[3]

In this paper we adopt the hybrid approach. Thus, at the core of our analysis is the set of subgame-perfect equilibria (SPE) of the base game. Implications of active negotiation by the players in the negotiation stage are expressed in the form of cooperative bargaining theory that selects among these equilibria.

In fact, at this point we do not need to assume any particular bargaining theory, because our main results require only that the bargaining theory satisfy one of the minimal collective-rationality conditions described in the next subsection. In Section 4 we refine further by providing the details of the hybrid model with a specific bargaining theory that entails application of the Nash bargaining solution (based on Miller and Watson, 2013a). The more selective model is used to produce additional results and numerical examples later in the paper.

Minimal collective-rationality assumption

In this subsection we define a notion of *internal consistency* that, in similar various forms, is incorporated into most major theories of collective rationality in repeated game theory, including contractual equilibrium as defined by Miller and Watson (2013b) and Watson, Miller, and Olsen (2020), as well as definitions of renegotiation-proofness including those of Bernheim and Ray (1989), Farrell and Maskin (1989), Bergin and MacLeod (1993), Ray (1994), and Goldlücke and Kranz (2013).

Consider the base game, where the strategy profile, denoted by s, specifies the players' actions in the transfer, production, and payment phases of each period, as a function of the history of play. Let h denote any history through the end of a period, and denote the initial (null) history by h^0. Let H be the space of all such histories.

[3]For the setting in which $C = N$ and $\rho = 1$, Miller and Watson (2013b) provides an equivalence result for a particular hybrid model (with the Nash bargaining solution) and an axiomatic equilibrium selection from a fully noncooperative counterpart (with a standard bargaining protocol such as alternating offers). The axioms imposed in the fully noncooperative model relate the outcome of negotiation in a period (such as whether a verbal agreement was reached) to the continuation value that the players coordinate to obtain in the continuation of the game from the transfer phase.

Definition: For any t-period history h and $(t+1)$-period history h', call h' an **immediate successor** of h if the t-period truncation of h' is exactly h.

In other words, h' is an immediate successor of h if it is the sequence of actions given by h followed by the sequence of actions in one additional period.

Let us introduce notation for describing the *continuation value* of any strategy profile s from any history h. Letting t denote the length of history h, the continuation value is the expected payoff vector that would be achieved in the subgame starting in period $t+1$ assuming play in the subgame is as s specifies. The continuation value does not include payoffs received earlier and is discounted to period $t+1$.

For a given strategy profile s and history h, let $v^s(h)$ denote the continuation value of s following history h, and let $V^s \equiv \{v^s(h) \mid h \in H\}$. Let $\overline{v}^s(h)$ denote the continuation value from the transfer phase of a period in which $\phi = 1$ (there is a renegotiation opportunity) and h is the history through the end of the previous period. Likewise, let $\tilde{v}^s(h)$ denote the continuation value from the transfer phase of a period in which $\phi = 0$ (there is no renegotiation opportunity) and h is the history through the end of the previous period. Note that $v^s(h) = \rho\overline{v}^s(h) + (1 - \rho)\tilde{v}^s(h)$.

We shall impose internal consistency in the following sense: In any period t, following any $(t-1)$-period history h, and in the event that there is an opportunity to negotiate in period t (i.e., $\phi_t = 1$), the coalition does not have the joint incentive to switch to an alternative equilibrium that is *viable* in relation to the current equilibrium. "Viable" means that the alternative equilibrium specifies behavior from $t+1$ that matches what the current equilibrium would specify following some history (not necessarily the same history), and so it amounts to continuation play that the players were already prepared to carry out. "Joint incentive to switch" means that every player in the coalition would be strictly better off switching to the alternative equilibrium. Here are the formal definitions:

Definition: Given a subgame-perfect equilibrium s of the base game and history h, call $w \in \mathbb{R}^n$ ***s*-viable from *h*** if there is a subgame-perfect equilibrium s' such that $\overline{v}^{s'}(h) = w$ and $v^{s'}(h') \in V^s$ for every $h' \in H$ that is an immediate successor of h.

Definition: Given a coalition C, call a base-game strategy profile s a ***coalition-C internally-consistent (C-IC) equilibrium*** if (1) it is a subgame-perfect equilibrium and (2) there does not exist a history $h \in H$ and vector $w \in \mathbb{R}^n$ such that w is s-viable from h, and $w_i > \overline{v}_i^s(h)$ for each player $i \in C$.

A C-IC equilibrium embodies our theory of the governing coalition, which is that the coalition's agreement on equilibrium play serves to coordinate not only the members of the coalition but also the outsiders. The ability to coordinate, subject to the individuals' incentive conditions, is the coalition's sole source of power. Condition 2 of the C-IC definition amounts to a modest exercise of this power: When given the opportunity to negotiate, the coalition will not settle for an equilibrium continuation if there is another viable equilibrium continuation that every coalition member would strictly prefer.[4]

For some applications, it may be unreasonable to expect that outsiders would make transfers to insiders in the transfer phase. For instance, an overarching norm may keep players from contemplating that an agreement made by the coalition at the beginning of a period could govern outsider transfers before any productive actions are taken in the period. In this case, we would limit attention to equilibria in which the players ignore transfers made by outsiders in the transfer phase, which, of course, implies that they will not be made in equilibrium.

Definition: Say that a strategy profile s exhibits ***neutral outsider transfers*** if the specified actions for each player do not depend on past actions taken by outsiders in the transfer phase of any period.

We next define a version of C-IC for the class of neutral outsider transfers.

Definition: Given a subgame-perfect equilibrium s, call $w \in \mathbb{R}^n$ ***s-viable with neutral outsider transfers*** if there is a subgame-perfect equilibrium s' that exhibits neutral outsider transfers such that $\overline{v}^{s'}(h) = w$ and $v^{s'}(h') \in V^s$ for every

[4]As noted above, some form of internal consistency is built into standard definitions of contractual equilibrium and renegotiation-proofness, and these concepts impose additional conditions. The definition of contractual equilibrium, in particular, adds more structure by incorporating an explicit theory of negotiation.

$h' \in H$ that is an immediate successor of h.

Definition: Given a coalition C, call a strategy profile s a ***coalition-C internally-consistent equilibrium with neutral outsider transfers (C-ICN*** if (1) s is a subgame-perfect equilibrium that exhibits neutral outsider transfers, and (2) there does not exist a history $h \in H$ and vector $w \in \mathbb{R}^n$ such that w is s-viable with neutral outsider transfers and $w_i > \overline{v}_i^s(h)$ for each player $i \in C$.

3 Motivating Outsiders and Insiders

The governing coalition has the dual concerns of maintaining cooperation within the coalition and motivating outsiders to behave as the coalition prefers. In both realms, the coalition would like to provide incentives using the promise of future rewards and punishments. However, these promises may be limited by the coalition's lack of commitment not to renegotiate. The prior literature has studied how the prospect of renegotiation complicates cooperation in standard repeated games, a setting that effectively assumes all players are in the governing coalition. In this section, we develop the main insights of our coalition analysis, showing that renegotiation also complicates the extent to which a coalition can motivate outsiders, and that these complications are in a sense more severe. Implications are generated for the form of agreements within the coalition and the form of (possibly implicit) agreements between coalition members and outsiders.

No-renegotiation benchmark

We begin with a benchmark result for the setting in which the coalition members negotiate in the first period only. This corresponds to the case of $\rho = 0$, and yet $\phi = 1$ in the first period, where the continuation value of an equilibrium s is $\overline{v}^s(h^0)$. Let V denote the set of subgame-perfect equilibrium payoff vectors in the base game. For any compact set $W \subset \mathbb{R}^n$, set $C \subset N$, and player $i \in N$, define

$$\zeta_i(W) \equiv \min\{w_i \mid w \in W\}$$

and

$$L^C(W) \equiv \max_{w \in W} \sum_{j \in C} w_j.$$

As is standard in game-theory notation, for any player $i \in N$, we let $-i \equiv N \setminus \{i\}$ denote the set of players other than i. For instance, a_{-i} is the vector of productive actions taken by the players other than i. Also, let $-C \equiv N \setminus C$ denote the set of outsiders.

Theorem 1. *Consider any base game and coalition $C \subset N$ and assume that V, the set of subgame-perfect equilibrium payoffs, is compact. In the case of $\rho = 0$, there is a C-IC equilibrium $\overline{s}$ such that $\sum_{i \in C} v_i^{\overline{s}}(h^0) = L^N(V) - \sum_{j \in -C} \zeta_j(V)$.*

That is, there is a C-IC equilibrium on the Pareto frontier of the set of SPE in which the coalition extracts all of the value, relative to the minimum that each outsider must get in an equilibrium. This is indeed the best outcome that the coalition could hope for, and therefore we might expect the coalition to coordinate in this fashion.

The Appendix contains a proof of Theorem 1 and proofs of the other theorems presented herein. The logic underlying Theorem 1 runs as follows: One can find a SPE s^* that maximizes the sum of all players' first-period continuation values. Also, for every player $i \in -C$, one can find an SPE s^i that minimizes player i's continuation value, over all SPE continuations, and also entails no transfers in the first-period transfer phase. Then we can construct an SPE in which, in the first period, every player $i \in -C$ transfers to a coalition member the difference between player i's continuation value of s and player i's continuation value of s^i, under threat of switching to s^i if player i fails to make the prescribed transfer. By construction, s^* achieves joint value stated in the theorem.

Insider continuation values

Returning to the general model with $\rho \in [0, 1]$, let us explore how the opportunity to renegotiate affects the coalition's equilibrium continuation values. The CIC condition and ability to make transfers leads to the following characterization result:

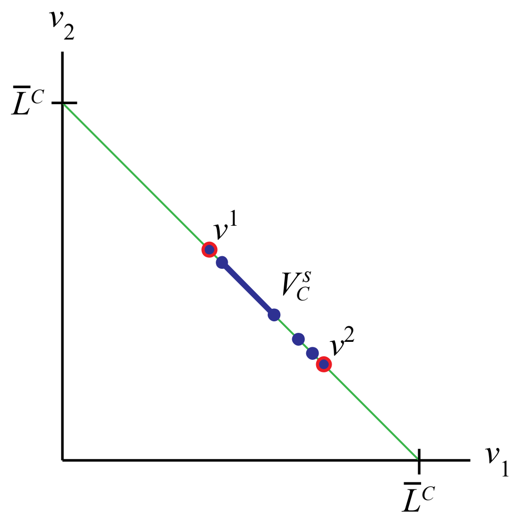

Figure 1: C-IC equilibrium coalition continuation values contingent on negotiation.

Theorem 2. *If s is a C-IC or C-ICN equilibrium, then there is a joint value $\overline{L}^C \in \mathbb{R}$ such that $\sum_{i\in C} \overline{v}_i^s(h) = \overline{L}^C$ for every $h \in H$.*

Figure 1 illustrates the result in the case of a two-player coalition $C = \{1, 2\}$. The thin green line of slope -1 gives the set of vectors with joint value (the sum of the insiders' individual values) of $\overline{L}^C$. The points shown in blue are the various equilibrium continuation-value vectors $\overline{v}_i^s(h)$ for $h \in H$. The set of these vectors is V_C^s, the projection of V^s to the coalition members' values. In this example, z^2 is the equilibrium continuation value least favorable to player 2, and z^1 is the equilibrium continuation value least favorable to player 1.

Theorem 2 follows from a straightforward characterization of achievable continuation values from the transfer phase of any given period, in relation to the set of achievable continuation values from the action phase. Intuitively, when negotiation takes place, the coalition members aim to maximize their joint value, and can divide this value arbitrarily by making transfers between themselves. Pareto inferior equilibrium continuation values can be improved upon in a way that obeys incentive conditions in the transfer phase. The repeated game is stationary, and so the same set of continuation values from the next period applies for the characterization of equilibrium values from any period.

Outsider continuation values

Next let us explore how the opportunity to renegotiate affects the coalition's handling of outsiders. The following result identifies a property of C-IC equilibria with interesting economic implications.

Theorem 3. *If s is a C-IC equilibrium then $\{\overline{v}^s_{-C}(h) \mid h \in H\}$ is a singleton.*

That is, in any C-IC equilibrium, the vector of outsiders' continuation values contingent on renegotiation is constant across all periods and histories. In other words, the outsiders' continuation values must reset at every instance of negotiation by the coalition. An important implication is that outsiders cannot be motivated by the promise of future rewards and punishments that would follow the next negotiation opportunity, because such rewards and punishments cannot be supported in a C-IC equilibrium.

For ρ close to 1, Theorem 3 establishes as a result of the coalition's collective rationality what the prior literature on international environmental agreements frequently assumed: that outsiders behave myopically from period to period. However, in settings with $\rho < 1$, there still is scope for rewards and punishments to take place in future periods. Further, even in the case of $\rho = 1$, outsiders may be induced to behave non-myopically across phases *within* a period.

The proof of Theorem 3 takes advantage of the ability of outsiders to make transfers to coalition members in the transfer phase. Essentially, the coalition asks each outsider to pay into the coalition an amount that pushes the outsider down to this player's minimum equilibrium continuation value, on the extortive threat to coordinate on this equilibrium continuation in the event the outsider does not pay. However, the conclusion of Theorem 3 also holds for CICN equilibria, which do not employ extortive transfers, under an additional mild condition on the shape of the payoff functions.

For any $a \in A$ and $i \in N$, let $\hat{u}_i(a) \equiv \max_{a'_i \in A_i} u_i(a'_i, a_{-i})$, assuming the maximum exists. Note that a sufficient condition for the maximum to exist is that A_i is compact and u_i is continuous in a_i.

Definition: Say that ***the production technology is well-behaved*** if, for every $i \in N$, A_i is a convex set, u_i is continuous and strictly concave, and $\hat{u}_i - u_i$

is convex.

It is common in the applied literature to make functional-form assumptions that satisfy the above condition. In an environmental application, for instance, player i's action a_i could be abatement, emissions reduction, investment, or the like. Concavity of u_i means that abatement exhibits decreasing returns. Convexity of $\hat{u}_i - u_i$ means that player i's deviation gain increases in the players' abatement levels; in particular, player i obtains a greater deviation gain when deviating from a large abatement level (to, say, zero) than with a smaller level.

Theorem 4. *If the production technology is well-behaved and s is a C-ICN equilibrium, then $\{\overline{v}^s_{-C}(h) \mid h \in H\}$ is a singleton.*

As before, the continuation values of the outsiders must reset at every instance of negotiation by the coalition, and for ρ close to 1, outsiders behave myopically from period to period (not necessarily myopically in each stage within a period). The proof of Theorem 4 is based on the following logic: Recall that $\overline{L}^C$ denotes the coalition's singular joint value contingent on negotiation. Consider a C-ICN equilibrium s, and suppose that $\overline{v}^s_{-C}(h) \neq \overline{v}^s_{-C}(h')$ for two histories h and h'. We can show that these values are achieved with different action profiles in the action phase of the current period. That is, some action profile a is chosen after h, whereas a' is chosen after h'. One can then find a continuation value w that is s-viable with neutral outsider transfers and is constructed by inducing play of the average action profile, $(a + a')/2$, in the current period. The construction is possible utilizing continuation values in V^s due to $\hat{u}_i - u_i$ being convex (that is, the incentive conditions in the action phase are weakly slackened). It turns out that w delivers for the coalition a strictly higher joint value than $\overline{L}^C$ owing to the strict concavity of u_i, which is a contradiction that renders our initial presumption false.

Equilibrium insider and outsider incentive arrangements

Let us explore how the coalition provides incentives for insiders and outsiders in a C-IC equilibrium. The following logic holds regardless of whether we assume neutral outsider transfers, and so we shall refer to "C-IC equilibrium" without specifying whether neutrality to outsider transfers is assumed.

First consider the nature of outsiders' incentives in the action phase of a given period. As Theorems 3 and 4 establish, the outsiders cannot face rewards and punishments that would occur after the next instance of negotiation. However, importantly, this does not imply myopic behavior in the productive phase, for two reasons. First, if $\rho < 1$ then negotiation may not occur in the following period, and in this case the outsiders' continuation values may depend on the history of play including the productive actions in the current period. Second, even if $\rho = 1$, outsiders may be motivated to play non-myopically in the action phase owing to end-of-period payments that are contingent on their productive actions.

In contrast, insiders are subject to rewards and punishments occurring after the next negotiation opportunity, because for a C-IC equilibrium s, generally $\{\overline{v}^s_C(h) \mid h \in H\}$ is not a singleton. That is, there are multiple (non-Pareto-ranked) vectors of equilibrium continuation values for coalition members. Coalition members may be motivated to take specific actions within a period by the promise of future-period, post-negotiation-opportunity rewards and punishments. Further, the actions motivated in this fashion are not restricted to those taken in the action phase. Insiders can also be motivated to make prescribed payments, including to outsiders, on the basis of future rewards and punishments, even for $\rho = 1$.

Thus emerges a broad theme: *In equilibrium, insiders are motivated by longer-term incentive arrangements, whereas outsiders are motivated by shorter-term incentive arrangements.* In contractual language, the coalition members form a long-term relational contract, and they have a (perhaps implicit) shorter-term relational contract with the outsiders. The latter contract specifies how payments from insiders to outsiders are conditioned on the outsiders' productive actions take earlier in the period. The coalition's long-term contract provides the incentives for insiders to make these payments. That is, if a coalition member deviates by not making a payment to a particular outsider, then this coalition member is punished by a shift in how the coalition coordinates its future behavior.

To explore the logic more precisely, and to see the trade-offs that the coalition resolves in a C-IC equilibrium, consider the setting of $\rho = 1$, so that negotiation is triggered in every period. Then the relational contracts with outsiders are effectively one-period contracts and they rely on enforcement entirely within each period. The relational contract between coalition members is a long-term contract, with enforcement taking place across periods.

Take any C-IC equilibrium s. Recall that V^s gives the set of continuation values supported by the equilibrium, and V_C^s is the projection on the space of values for the coalition members. When the coalition members negotiate at the beginning of a period, they consider how best to prescribe behavior in the transfer, production, and payment phases of the current period. Correspondingly, they specify how to coordinate on a continuation value in V_C^s from the beginning of the next period, as a function of play in the current period. Further, the latter must ensure incentive-compatibility of the former. In other words, the manner of coordination on the next-period continuation value provides incentives for the players to behave as prescribed in the current period (self enforcement). Multiple, different values in V_C^s can be used to reward and punish insiders.

For example, consider the case of $C = \{1, 2\}$, as shown in Figure 1. Coordinating on continuation value z^1 from the next period would impart the maximum punishment on player 1, for any deviation in the current period. Likewise, coordinating on continuation value z^2 from the next period would maximally punish player 2. Coordinating on an intermediate value would reward them both for behaving as prescribed in the current period. Thus, the *span* of the set V_C^s, defined as $z_1^2 - z_1^1 = z_2^1 - z_2^2$, gives the capacity for rewards and punishments that is available, and utilized, in equilibrium s.

What insider actions is the limited span used to motivate in equilibrium? Consider the following backward-induction heuristic, which will turn out to be justified by the analysis shown in the Appendix: The span of continuation values from the next period can be viewed as giving the coalition members an incentive to make payments at the end of the current period, and in turn these payments are used to provide incentives in the action phase of the current period.

Therefore, one can think of the span as a total bonus, divided between providing production-phase incentives for coalition members and providing production-phase incentives for outsiders. The coalition's long-term contract optimally allocates these shares. The tradeoff is simple. We can calculate the marginal value to the coalition that an incremental bonus would achieve if spent to motivate outsiders to take productive actions in the coalition's interest. Likewise, we can calculate the marginal value to the coalition that the same incremental bonus would achieve if spent to motivate insiders. The coalition's equilibrium contract optimally allocates the bonus to the group (insiders or outsiders) that gives the higher return. In a setting with a continuous space of productive actions, an interior solution equates the marginal returns of outsider- and insider-channeled bonuses.

For some parameter values, the equilibrium contract may allocate most of the bonus to provide incentives to outsiders. In this case, the short-term contracts between insiders and outsiders has the outsiders taking costly productive actions (that provide benefit to the coalition) in each period, rewarded by payments from coalition members at the end of the period. The long-term contract between coalition members is a self-enforced agreement to make the payments to outsiders that the short-term contracts prescribe.

4 Coalition-Directed Contractual Equilibrium

In this section we introduce *coalition-directed contractual equilibrium (CDCE)* to derive a more specific equilibrium selection. Contractual equilibrium incorporates the weighted Nash bargaining solution as the bargaining protocol in the negotiation phase. If the coalition that directs the equilibrium includes all the players and negotiation opportunities arise every period, then CDCE is equivalent to contractual equilibrium as defined in Miller and Watson (2013a).

Coalition-directed contractual equilibrium describes the joint behavior of insiders in the negotiation and transfer phases and the non-cooperative behavior of all players in the production and payment phases in each period. For simplicity, in this paper we employ a recursive definition of CDCE, and restrict attention to pure strategies.[5] In the negotiation phase, if a negotiation opportu-

[5]Watson et al. (2020), which employs a related model with cooperative bargaining, defines contractual equilibrium using both a recursive approach and a strategic approach, and discusses the relationship between the two approaches. Miller and Watson (2013b) shows how axiomatic

nity arises then the insiders engage in weighted Nash bargaining with transfers, where the vector of *bargaining weights* is $\pi = (\pi_i)_{i \in C} \in \Delta C$. Let $\hat{W} \subset \mathbb{R}^n$ be the set of continuation values available from the start of the action phase. If the insiders fail to agree, then their continuation play will yield some particular *disagreement value* $\underline{w} \in \hat{W}$. In equilibrium, they agree to make an immediate transfer m and then continue with some value $w \in \hat{W}$, so as to obtain their Nash-bargaining payoff vector $w_C + m_C = \underline{w}_C + \pi_C\big(L^C(\hat{W}) - \sum_{i \in C} \underline{w}_i\big)$. In the transfer phase, transfers can occur only if a negotiation opportunity arose and agreement was reached. There are no transfers between insiders and outsiders, so CDCE is neutral to outsider transfers. Any deviation in the transfer phase triggers a disagreement. Of course, in the production and payment phases each player—insider or outsider—must prefer not to unilaterally deviate. We assume that the production technology is well-behaved.

A formal definition of CDCE is reserved to the Appendix. Here we employ some backward induction within the period and the fact that CDCE is C-ICN to simplify the description of CDCE, without any loss of generality.

First consider the payment phase. Let $W \subset \mathbb{R}^n$ be the set of continuation values available from the start of the next period. Let $p^i \in \mathbb{T}^n(i)$ and $w \in W$ be the payment each player i is supposed to make and the continuation value the players are supposed to receive, given whatever has happened up to the start of the payment phase. Player i's best deviation in the payment phase is to choose the payment $(0, \ldots, 0)$ instead of p^i, for an immediate gain of $-p^i_i \geq 0$. Enforcing p^i thus requires that $-(1-\delta)p^i_i \leq \delta\big(w_i - \zeta_i(W)\big)$, since $\zeta_i(W)$ is the worst continuation payoff that can be imposed on player i. We restrict attention to equilibria in which any unilateral deviation in the payment phase is punished in this manner, without loss of generality. Adding up across players yields the requirement that

$$-(1-\delta)\sum_i p^i_i \leq \delta\Big(L^N(W) - \sum_i \zeta_i(W)\Big). \tag{1}$$

That is, the total amount that players can be made to pay cannot be more than

equilibrium selection in a fully non-cooperative model provides foundations for models with cooperative bargaining.

the difference between the maximum joint payoff and the sum of individual players' minimum payoffs. The requirement holds with equality if $\sum_i w_i = L^N(W)$.

Next consider the action phase. Let a be the action profile that is supposed to be chosen given whatever has happened up to the start of the action phase, and let p be the payment that is supposed to be made in the payment phase if a was chosen in the action phase. Let player i's best deviation in the action phase yields an immediate payoff of $\hat{u}_i(a)$, so the gain from deviating is $\hat{u}_i(a) - u_i(a)$. To enforce a, the equilibrium needs to punish such a deviation by some combination of reducing player i's net payment and player i's continuation value. One way to impose the maximal punishment after player i's production-phase deviation is for all other players to pay zero to player i, rather than $\sum_j p_i^j$, for player i to pay zero to all other players, and for player i's continuation value to be set to ζ_i; the magnitude of this punishment is $(1-\delta)\sum_{j \in N} p_j^i + \delta\big(w_i - \zeta_i(W)\big)$ (note that p_i^i, which is non-positive, is included in the sum, because not having to pay it reduces the magnitude of the punishment). However, another way to impose a punishment of the same magnitude is for all other players to pay zero to player i, player i to pay $\delta\big(w_i - \zeta_i(W)\big)$ in total to the other players, and for player i's continuation value to be set to w_i. Player i's payment in this case is enforced by the threat of having her continuation value set to $\zeta_i(W)$ if she fails to pay. To deter player i's action phase deviation, we must have $(1-\delta)\hat{u}_i(a) \leq (1-\delta)\sum_{j \in N} p_j^i + \delta\big(w_i - \zeta_i(W)\big)$. Adding across players yields the requirement

$$
\begin{aligned}
(1-\delta)\sum_i \hat{u}_i &\leq (1-\delta)\sum_i \sum_j p_j^i + \delta\Big(L^N(W) - \sum_i \zeta_i(W)\Big) \\
&\leq \delta\Big(L^N(W) - \sum_i \zeta_i(W)\Big),
\end{aligned}
\tag{2}
$$

where the second inequality holds because $\sum_i \sum_j p_j^i = \sum_j \sum_i p_j^i \leq 0$.

Equation (2) shows that even though there is a payment phase at the end of each period, the aggregate enforcement power (to discourage deviations from the desired productive action profile) is nonetheless limited by the punishments available in the continuation game that starts in the following period. In addition, Equation (1)—and the fact that it can be made to hold with equality—shows that the players can use payments alone to enforce their desired productive action profile, and use changes in continuation values only to enforce those payments. Goldlücke and Kranz (2012) showed this holds generally for subgame perfect equilibria in repeated games with transfers.

Now consider the negotiation and transfer phases. If there is no negotiation opportunity ($\phi = 0$), then there is nothing for the players to do. If there is a negotiation opportunity, then these two phases are jointly resolved by the weighted Nash bargaining solution. Under agreement the insiders make transfers m_C and the players receive continuation payoff vector $w \in \hat{W}$ starting in the action phase, so coalition members receive the payoff $w_C + m_C = \underline{w}_C + \pi_C\big(L^C(\hat{W}) - \sum_{i \in C} \underline{w}_i\big)$ from the start of the period. If there is disagreement or a deviation in the transfer phase, then the players receive continuation payoff vector $\underline{w} \in \hat{W}$ from the start of the action phase.

To formally define CDCE under these simplifcations, let $\tau : A \to \bigtimes_{i \in N} \mathbb{R}^n_0(i)$ specify what the strategy profile calls for in the payment phase as a function of the production-phase outcome, and let $y : A \times \bigtimes_{i \in N} \mathbb{R}^n_0(i) \to W$ specify the continuation value as a function of the production-phase action profile and realized payments. Then in equilibrium (a, τ) must be a subgame perfect equilibrium in the reduced game with payoff function

$$g(a, \tau, y) \equiv (1 - \delta)\Big(u(a) + \sum_i \tau_i(a)\Big) + \delta y(a, \tau(a)). \tag{3}$$

If this is the case, we say that y *enforces* a or (a, τ) *relative to* W.

Given W, let $D(W)$ be the set of equilibrium values from the action phase of the current period. That is,

$$D(W) \equiv \left\{ w \in \mathbb{R}^n \;\middle|\; \begin{array}{l} \exists\, a \in A,\ \tau : A \to \bigtimes_{i \in N} \mathbb{R}^n_0(i), \text{ and } y : A \times \bigtimes_{i \in N} \mathbb{R}^n_0(i) \to W \\ \text{s.t. } y \text{ enforces } (a, \tau) \text{ relative to } W \text{ and } w = g(a, \tau, y). \end{array} \right\} \tag{4}$$

Since any $\underline{w} \in D(W)$ can be the disagreement value, the set of achievable continuation values from the negotiation phase is

$$B^C(W) \equiv \left\{ \left(\underline{w}_C + \pi_C \Big(L^C(W) - \sum_{i \in C} \underline{w}_i \Big),\ w_{-C} \right) \,\middle|\, \underline{w} \in D(W),\ (w_C, w_{-C}) \in \Lambda^C(W) \right\}. \tag{5}$$

Definition 1. *The* contractual equilibrium value (CEV) set for coalition C, *denoted by* W^{C*}, *is a fixed point of* $B^C(D(\cdot))$ *that is maximal in the sense of having the highest level* $L^C(W)$.

We first prove some preliminary results that simplify the problem of computing the CEV set for a coalition.

We say that $y : A \times \bigtimes_{i \in N} \mathbb{R}^n_0(i) \to W$ *Nash-enforces* (a, τ) *relative to* W if (a, τ) is a Nash equilibrium in the reduced game comprising the production and payment phases, where the reduced-game payoff function is $g(a, \tau(a), y)$. We show that the structure of payments implies that it is without loss of generality to focus on Nash enforcement rather than subgame perfect enforcement in the reduced game, as if players announce and commit to their payment-phase strategies at the same time as they choose their production-phase actions.

Lemma 1. *Given* $W \subset \mathbb{R}^n$,

$$D(W) = \left\{ w \in \mathbb{R}^n \,\middle|\, \begin{array}{l} \exists a \in A,\ \tau : A \to \bigtimes_{i \in N} \mathbb{R}^n_0(i),\ \text{and } y : A \times \bigtimes_{i \in N} \mathbb{R}^n_0(i) \to W \\ \text{s.t. } y \text{ Nash-enforces } (a, \tau) \text{ relative to } W \text{ and } w = g(a, \tau, y). \end{array} \right\} \tag{6}$$

Proof. Since subgame perfect enforcement implies Nash enforcement, it suffices to show that any w that can be Nash-enforced relative to W can also be (subgame-perfect) enforced relative to W. So suppose that y Nash-enforces (a, τ) relative to W and $w = g(a, \tau, y)$. We construct an alternative $\hat{y}$ and $\hat{\tau}$ such that $\hat{y}$ (subgame-perfect) enforces $(a, \hat{\tau})$. First, we construct $\hat{y}$ as follows. Any unilateral deviation by player i from (a, τ) can be punished at least as harshly by continuing with $\zeta_i(W)$ rather than what is specified by y, so construct a provisional $\hat{y}$ by modifying y to implement this policy. Moreover, since multilateral deviations are irrelevant to Nash enforcement, further modify $\hat{y}$ to specify that the continuation value is simply $y(a)$ following all multilateral deviations.

Next, we construct $\hat{\tau}$ as follows. Each player $j \in N$ can specify zero payments (to all other players) for any $a' \neq a$ so as to punish any production-phase deviation at least as harshly as specified by τ, so construct $\hat{\tau}$ by modifying τ to implement this policy.[6] Under $(a, \hat{\tau}, \hat{y})$, any best deviation by player i yields a payoff in the reduced game of $(1-\delta)\hat{u}_i(a) + \delta\zeta_i(W)$, since under this punishment policy a best deviation is to maximize her payoff in both the action phase (obtaining $\hat{u}_i(a)$) and the payment phase (e.g., by not choosing a deviant $\hat{\tau}$ because the equilibrium $\hat{\tau}$ already specifies that she should not make any payments after deviating in the action phase).

With these modifications, $\hat{y}$ now (subgame-perfect) enforces $(a, \hat{\tau})$ relative to W. First, we have already shown above that no unilateral deviation in the action phase is profitable. Second, following any (unilateral or multilateral) deviation $a' \neq a$ in the action phase, all players make no payments and then receive a continuation value that does not depend on their payments; this is a Nash equilibrium in the payment-phase subgame of the reduced game after a' is chosen in the action phase. Finally, if there is no deviation in the action phase, the equilibrium-path payment profile $\hat{\tau}(a) = \tau(a)$ is a Nash equilibrium in the in the subsequent payment-phase subgame, since any unilateral deviator suffers at least as harsh a punishment as in the original (Nash-enforced) (a, τ, y). Therefore, $(a, \hat{\tau})$ is enforced by $\hat{y}$, and by construction $g(a, \hat{\tau}, \hat{y}) = w$. It follows that $w \in D(w)$. □

Let tri $W \equiv \left\{ w \in \mathbb{R}^n \;\middle|\; w_i \geq \zeta_i(W) \text{ for all } i \in N,\ \sum_{i \in N} w_i \leq L(W) \right\}$.

Lemma 2. *$D(W) = D(\text{tri } W)$.*

Corollary 1. *W^{C*} has n extreme points. For each $i \in N$, there is a corresponding extreme point w^i such that $w^i_j = \zeta_j(W^{C*})$ for all $j \neq i$ and $w^i_i = L^C(W^{C*}) -$*

[6] Because there is perfect monitoring, the proof of the lemma is simplified by zeroing out all payments following a production-phase deviation. Alternatively, the players could construct more targeted punishments by zeroing out only their payments to the deviator, preserving what they would have paid each other under τ. While practical, such an approach would not suffice to prove the lemma, because τ could involve promising off-path payments that are not sequentially rational.

$\sum_{j \neq i} \zeta_j(W^{C*})$.

To simplify notation for the dynamic enforcement constraints, we introduce the following notation:

- $U^C(a) \equiv \sum_{i \in C} u_i(a)$ and $\hat{U}^C(a) \equiv \sum_{i \in C} \hat{u}_i(a)$,
- $U^{-C}(a) \equiv \sum_{i \notin C} u_i(a)$ and $\hat{U}^{-C}(a) \equiv \sum_{i \notin C} \hat{u}_i(a)$,
- $Z_C(W) \equiv \sum_{i \in C} \zeta_i(W)$.

Theorem 5. *The CEV set W^{C*} for coalition C has the following properties:*

1. *First, d^* is the largest fixed point of $\Gamma^C : \mathbb{R}_+ \to \mathbb{R}_+$, where $\Gamma^C(d) \equiv \sum_{i \in C} \gamma_i(d)$ and*

$$\begin{aligned} \gamma_i(d) \equiv \max_{a \in A} \, & \pi_i\big(U(a) - \hat{U}^{-C}(a)\big) - \hat{u}_i(a) \\ s.t. \; & \hat{U}(a) - U(a) \leq \frac{\delta}{1-\delta} d; \end{aligned} \tag{7}$$

2. *Second,*

$$\begin{aligned} L^C(W^{C*}) = \max_{a \in A} \, & U(a) - \hat{U}^{-C}(a) \\ s.t. \; & \hat{U}(a) - U(a) \leq \frac{\delta}{1-\delta} d^*; \end{aligned} \tag{8}$$

3. *Third, $d^* \equiv L^C(W^{C*}) - Z_C(W^{C*})$.*

Proof. Follows from the main theorem of Miller and Watson (2025) and Theorem 4. Let $P_{-C} = -\sum_{i \in C} \tau_i^i(a)$ be the aggregate amount that the coalition pays the outsiders on path, when the equilibrium action profile a is played in the action phase. Part 1 of the theorem shows how to compute the disagreement point to punish each player i, using a convex optimization problem with two simple constraints.

Before eliminating P_{-C}, we had

$$\gamma_i(d) \equiv \max_{a \in A; P_{-C} \in \mathbb{R}_+} \pi_i\big(U^C(a) - P_{-C}\big) - \hat{u}_i(a)$$
$$\text{s.t. } \hat{U}^C(a) - \big(U^C(a) - P_{-C}\big) \leq \frac{\delta}{1-\delta} d, \tag{9}$$
$$\hat{U}^{-C}(a) - \big(U^{-C}(a) + P_{-C}\big) \leq 0;$$

and

$$L^C(W^{C*}) = \max_{a \in A; P_{-C} \in \mathbb{R}_+} U^C(a) - P_{-C}$$
$$\text{s.t. } \hat{U}^C(a) - \big(U^C(a) - P_{-C}\big) \leq \frac{\delta}{1-\delta} d^*, \tag{10}$$
$$\hat{U}^{-C}(a) - \big(U^{-C}(a) + P_{-C}\big) \leq 0;$$

In each case the outsiders' aggregate enforcement constraint must bind: suppose not; then P_{-C} can be reduced without violating that constraint, while the dynamic enforcement constraint is relaxed and the objective is improved. Accordingly, substituting for P_{-C} yields Equations (7) and (8).

□

5 Numerical Example: Climate Change Treaties

To model coalitional negotiation of climate change treaties, consider the following game.

- There are n nations $i \in \{1, \dots, n\} \equiv N$, each with a fixed and publicly known size $\theta_i \in \mathbb{R}_+$.

- In the action phase, each nation i's action is its per capita effort toward emissions abatement, $a_i \in \mathbb{R}_+$.

- Each nation i's production-phase payoff is $u_i(a) = \theta_i \sum_j \theta_j a_j - \psi_i \theta_i a_i^2$. That is, nation i's payoff includes its citizens' gains from global abatement effort $\theta_i \sum_j \theta_j a_j$ and its citizens' costs of abatement $-\psi_i \theta_i a_i^2$, where $\psi_i > 0$ is its fixed, idiosyncratic, and publicly known cost parameter.

- FOR NOW: We restrict attention to the case in which $\rho = 1$ so there is a renegotiation opportunity every period.

Modeling the benefits of mitigation effort as linear and the costs as quadratic follows Nordhaus (AER 2015) and Barrett (various), and is not particularly important to this exercise. A more important limitation is that we model this as a repeated game, wherein the costs and benefits of mitigation effort accrue only within a single period. Allowing for accumulation of both greenhouse gases and mitigation capital would dramatically complicate the analysis, requiring even further new generalization of contractual equilibrium to allow for endogenous state variables.[7]

5.1 Analysis

Preliminary calculations:

$$\frac{du_i(a)}{da_i} = \theta_i^2 - 2\psi_i\theta_i a_i \tag{11}$$

$$\hat{a}_i \equiv \arg\max_{a_i \in A_i} u_i(a) = \frac{\theta_i}{2\psi_i} \text{ for all } a_{-i} \tag{12}$$

$$\hat{u}_i(a) = \theta_i \sum_{j \neq i} \theta_j a_j + \frac{\theta_i^3}{4\psi_i} \tag{13}$$

$$\hat{u}_i(a) - u_i(a) = \left(\frac{\theta_i^3}{2\psi_i} - \frac{\theta_i^3}{4\psi_i} - \theta_i^2 a_i + \psi_i\theta_i a_i^2\right) = \frac{\psi_i\theta_i}{4}\left(2a_i - \frac{\theta_i}{\psi_i}\right)^2 \tag{14}$$

$$\hat{U}^C(a) - U^C(a) = \sum_{i \in C} \frac{\psi_i\theta_i}{4}\left(2a_i - \frac{\theta_i}{\psi_i}\right)^2 \tag{15}$$

$$P_O^* = \hat{U}^O(a) - U^O(a) = \sum_{i \in O} \frac{\psi_i\theta_i}{4}\left(2a_i - \frac{\theta_i}{\psi_i}\right)^2 \tag{16}$$

[7] Watson et al. (2020) defined contractual equilibrium for a stationary environment in which parties could write a long-term externally-enforced contract, which constituted an endogenous payoff-relevant state variable, but showed that contractual equilibria involved stationary behavior on the equilibrium path. Equilibrium-path behavior need not be stationary if the model accounted for the accumulation of greenhouse gases and mitigation capital.

So on path we solve

$$
\begin{aligned}
L^C(W^{C*}) = \max_{a \in \mathbb{R}^n_+} \; & \sum_{i \in C} \theta_i \sum_{j \in N} \theta_j a_j - \sum_{i \in C} \psi_i \theta_i a_i^2 - \sum_{i \in O} \frac{\psi_i \theta_i}{4} \Big(2a_i - \frac{\theta_i}{\psi_i} \Big)^2 \\
\text{s.t.} \; & \sum_{i \in N} \frac{\psi_i \theta_i}{4} \Big(2a_i - \frac{\theta_i}{\psi_i} \Big)^2 \leq \frac{\delta}{1-\delta} d;
\end{aligned}
\tag{17}
$$

We are mainly interested in situations in which the CEV set features non-zero effort but falls short of the first best, so cases in which the dynamic enforcement constraint binds. We have the following first order conditions for a constrained non-zero solution, with multiplier $\mu^C > 0$:[8]

$$
\theta_i \sum_{j \in C} \theta_j - 2\psi_i \theta_i a_i - \mu^C \psi_i \theta_i \Big(2a_i - \frac{\theta_i}{\psi_i} \Big) = 0 \text{ for all } i \in C, \tag{18}
$$

$$
\theta_i \sum_{j \in C} \theta_j - (1 + \mu^C) \psi_i \theta_i \Big(2a_i - \frac{\theta_i}{\psi_i} \Big) = 0 \text{ for all } i \in O. \tag{19}
$$

These evaluate to

$$
a^*_{i \in C} = \frac{\sum_{j \in C} \theta_j + \mu^C \theta_i}{2\psi_i (1 + \mu^C)}, \tag{20}
$$

$$
a^*_{i \in O} = \frac{\sum_{j \in C} \theta_j + (1 + \mu^C) \theta_i}{2\psi_i (1 + \mu^C)}. \tag{21}
$$

Evidently an outsider is asked to exert more effort than an otherwise equivalent insider. On the other hand, for any outsider i, joining the coalition means exerting the same effort:

$$
\frac{\sum_{j \in C} \theta_j + (1 + \mu^C) \theta_i}{2\psi_i (1 + \mu^C)} = \frac{\sum_{j \in C+i} \theta_j + \mu^C \theta_i}{2\psi_i (1 + \mu^C)}
$$

since μ^C does not depend explicitly on the coalition membership, only on the efforts; coalition members' efforts, in turn, are increasing in the coalition membership except to the extent that the dynamic enforcement constraint is affected.

[8] If there is an unconstrained solution, then $\mu^C = 0$ and the constraints do not bind. If there is a solution with $a_j = 0$ for some j, then the first order condition for that a_j need only satisfy that the left-hand side is less than or equal to zero.

For each nation $i \in C$ to be punished under disagreement, we solve

$$
\begin{aligned}
\psi_i(d) \equiv \max_{a \in \mathbb{R}^n_+} \pi_i &\left(\sum_{j \in C} \theta_j \sum_{k \in N} \theta_k a_k - \sum_{j \in C} \psi_j \theta_j a_j^2 - \sum_{j \in O} \frac{\psi_j \theta_j}{4} \Big(2a_j - \frac{\theta_j}{\psi_j} \Big)^2 \right) \\
&- \left(\theta_i \sum_{j \neq i} \theta_j a_j + \frac{\theta_i^3}{4\psi_i} \right) \qquad (22) \\
\text{s.t. } &\sum_{j \in N} \frac{\psi_j \theta_j}{4} \Big(2a_j - \frac{\theta_j}{\psi_j} \Big)^2 \leq \frac{\delta}{1-\delta} d;
\end{aligned}
$$

We have the following first order conditions for a constrained non-zero solution, with multiplier $\lambda_i^C > 0$:

$$\pi_i \Big(\theta_i \sum_{k \in C} \theta_k - 2\psi_i \theta_i a_i \Big) - \lambda_i^C \psi_i \theta_i \Big(2a_i - \frac{\theta_i}{\psi_i} \Big) = 0 \text{ for } i, \qquad (23)$$

$$\pi_i \Big(\theta_j \sum_{k \in C} \theta_k - 2\psi_j \theta_j a_j \Big) - \theta_i \theta_j - \lambda_i^C \psi_j \theta_j \Big(2a_j - \frac{\theta_j}{\psi_j} \Big) = 0 \text{ for all } j \in C \setminus i, \qquad (24)$$

$$\pi_i \theta_j \sum_{k \in C} \theta_k - \theta_i \theta_j - (\pi_i + \lambda_i^C) \psi_j \theta_j \Big(2a_j - \frac{\theta_j}{\psi_j} \Big) = 0 \text{ for all } j \in O. \qquad (25)$$

These evaluate to

$$\tilde{a}_i^i = \frac{\pi_i \sum_{k \in C} \theta_k + \lambda_i^C \theta_i}{2\psi_i \big(\pi_i + \lambda_i^C \big)}, \qquad (26)$$

$$\tilde{a}_{j \in C \setminus i}^i = \frac{\pi_i \sum_{k \in C} \theta_k - \theta_i + \lambda_i^C \theta_j}{2\psi_j \big(\pi_i + \lambda_i^C \big)}, \qquad (27)$$

$$\tilde{a}_{j \in O}^i = \frac{\pi_i \sum_{k \in C} \theta_k - \theta_i + (\pi_i + \lambda_i^C) \theta_j}{2\psi_j \big(\pi_i + \lambda_i^C \big)}. \qquad (28)$$

Evidently the nation being punished exerts more effort than its coalition partners. Outsiders are asked to exert more effort than the coalition partners not being punished.

5.1.1 Example: Three symmetric nations

To gain intuition, consider the case in which $n = 3$, $c = 2$, $\theta_i = \psi_i = 1$ for all $i \in N$ and $\pi_i = 1/2$ for all $i \in C$. Due to symmetry, λ_i^C does not depend on i, so let $\lambda_i^C = \lambda^C$ for all $i \in C$. Then we have

$$a^*_{i \in C} = \frac{2 + \mu^C}{2 + 2\mu^C}, \tag{29}$$

$$a^*_{i \in O} = \frac{3 + \mu^C}{2 + 2\mu^C}, \tag{30}$$

and

$$\tilde{a}_i^i = \frac{1 + \lambda^C}{1 + 2\lambda^C}, \tag{31}$$

$$\tilde{a}^i_{j \in C \setminus i} = \frac{\lambda^C}{1 + c\lambda_i^C}, \tag{32}$$

$$\tilde{a}^i_{j \in O} = \frac{1}{2}. \tag{33}$$

Evaluating and simplifying the binding dynamic enforcement constraint when punishing nation i, and solving for λ^C, we find that

$$\lambda^C(d) = \frac{1}{2}\left(\sqrt{\frac{1 - \delta}{2\delta d}} - 1\right), \tag{34}$$

which requires $1 - \delta > 2\delta d$. When instead $1 - \delta < 2\delta d$, the dynamic enforcement constraint does not bind, so $\tilde{a}_i^i = 1$ and $\tilde{a}^i_{j \in C \setminus i} = 0$.

Then the span d^* is the largest fixed point of $\Gamma(d)$, which in this example evaluates to $d^* = \delta - \delta^2 + \sqrt{(1 - \delta)^2 \delta^2}$ when the dynamic enforcement constraint binds, and $d^* = \frac{1}{2}$ when it does not bind.

5.2 Joining the coalition

Suppose that the current coalition C invites outsider i to join, and the coalition extracts an initiation fee that makes i indifferent over whether to join. By rejecting the invitation, the outsider continues to earn a constant per-period payoff of

$$w_i^{C*} = u_i(a^{C*}) + P_i(a^{C*}) = \frac{\gamma_i\theta_i}{4}\Big(2a_i^{C*} - \frac{\theta_i}{\gamma_i}\Big)^2 + \theta_i \sum_{j\in N} \theta_j a_j^{C*} - \gamma_j\theta_i(a_i^{C*})^2. \quad (35)$$

For the coalition to want to issue the invitation, the coalition's per-period joint value must increase by at least this much if i joins, creating coalition $C + i$. Indeed, the coalition's joint value increases by

$$\begin{aligned} &\sum_{j\in C} \theta_j \Bigg(\sum_{k\in N} \theta_k \big(a_k^{(C+i)*} - a_k^{C*}\big) - \gamma_j\big((a_j^{(C+i)*})^2 - (a_j^{C*})^2\big)\Bigg) \\ &\quad + \theta_i\Bigg(\sum_{k\in N} \theta_k a_k^{(C+i)*} - \gamma_i (a_i^{(C+i)*})^2\Bigg) \\ &\quad - \sum_{j\in O-i} \frac{\gamma_j\theta_j}{4}\Bigg(\Big(2a_j^{(C+i)*} - \frac{\theta_j}{\gamma_j}\Big)^2 - \Big(2a_j^{C*} - \frac{\theta_j}{\gamma_j}\Big)^2\Bigg) + \frac{\gamma_i\theta_i}{4}\Big(2a_i^{C*} - \frac{\theta_i}{\gamma_i}\Big)^2 \end{aligned} \quad (36)$$

NOTE: Look first at $n-1$ coalition adding a member to form the grand coalition. Could also look at simplest $n = 3$ case.

Can we prove that grand coalition forms?

6 Endogenizing coalition expansion

Future work

References

James Bergin and W. Bentley MacLeod. Efficiency and renegotiation in repeated games. *Journal of Economic Theory*, 61(1):42–73, October 1993.

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

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

Susanne Goldlücke and Sebastian Kranz. Infinitely repeated games with public monitoring and monetary transfers. *Journal of Economic Theory*, 147(3):1191–1221, May 2012.

Susanne Goldlücke and Sebastian Kranz. Renegotiation-proof relational contracts. *Games and Economic Behavior*, 80:157–178, July 2013.

David A. Miller and Joel Watson. A theory of disagreement in repeated games with bargaining. *Econometrica*, 81(6):2303–2350, November 2013a.

David A. Miller and Joel Watson. A theory of disagreement in repeated games with renegotiation. *Econometrica*, 81(6):2303–2350, November 2013b.

Debraj Ray. Internally renegotiation-proof equilibrium sets: Limit behavior with low discounting. *Games and Economic Behavior*, 6(1):162–177, January 1994.

Joel Watson, David A. Miller, and Trond E. Olsen. Relational Contracting, Negotiation, and External Enforcement. *American Economic Review*, 110(7): 2153–2197, July 2020.