A Theory of Disagreement in Repeated Games with Renegotiation

David A. Miller and Joel Watson*
University of California, San Diego

July 10, 2010

Abstract

This paper develops the concept of *contractual equilibrium* for repeated games with transferable utility, whereby the players negotiate cooperatively over their continuation strategies at the start of each period. Players may disagree in the negotiation phase, and continuation play may be suboptimal under disagreement. Under agreement, play is jointly optimal in the continuation game, and the players split the surplus (according to fixed bargaining weights) relative to what they would have attained under disagreement. Contractual equilibrium outcomes also arise from subgame perfect equilibria in a class of models with noncooperative bargaining, under some assumptions on the endogenous meaning of cheap-talk messages. Contractual equilibria exist for all discount factors, and for any given discount factor all contractual equilibria attain the same aggregate utility. Patient players attain efficiency under simple sufficient conditions; necessary and sufficient conditions are also provided. The allocation of bargaining power can dramatically affect aggregate utility. The theory extends naturally to games with more than two players, imperfect public monitoring, and heterogeneous discount factors.

JEL: C71 , C72 , C73 , C78

*We thank seminar participants at Collegio Carlo Alberto, Columbia, Florida International, the International Food Policy Research Institute, Penn, Toronto, UCLA, UCSD, UC Davis, USC, USC Marshall School of Business, Western Ontario, the 2009 NBER Organizational Economics Working Group Meeting (especially Sylvain Chassang), the Third World Congress of the Game Theory Society, the 2008 Southwest Economic Theory Conference, and the 2008 Stony Brook Workshop on Recent Advances in Repeated Games, for valuable comments and suggestions. Jacob Johnson and Jong-Myun Moon provided excellent research assistance. Both authors: Dept. of Economics, 9500 Gilman Dr. #0508, La Jolla, CA 92093–0508. Miller thanks Yale University and the Cowles Foundation for Research in Economics for hospitality and financial support, and Partha Dasgupta for inspiration to pursue this topic. Watson thanks the National Science Foundation (SES-0095207) and the National Oceanic and Atmospheric Administration (Southwest Fisheries Science Center), and the Cowles Foundation at Yale for financial support and hospitality.

Contents

1 Introduction

Many relationships (such as business partnerships, employment relationships, and buyer-supplier relationships) are ongoing and governed at least in part by self-enforced relational contracts. These long-term relationships are often subject to renegotiation, in a way anticipated by the agents involved. The renegotiation process lends each party some bargaining power, which he or she exercises to influence the terms of continued interaction. In the event that bargaining results in disagreement during a particular period of time, the parties may still continue interacting.

We seek to understand the outcomes that can arise when players in a long-run relationship renegotiate periodically, where players have bargaining power, and where they continue to interact whether or not they reach agreement when renegotiating. To this end, we introduce the concept of *contractual equilibrium* to describe contracting and renegotiation in repeated games with transfers. Our main message is that whether and how players can agree to sustain cooperation should be sensitive to how they would behave if they were to disagree.

The key element of our theory is an explicit account of negotiation activity, which occurs cooperatively in each period of the game, prior to the noncooperative stage-game actions. We describe behavior in the negotiation phase according to the cooperative Nash (1950) bargaining solution, where the players use transfers to split the available surplus relative to the disagreement outcome, according to fixed bargaining weights. If they disagree, they may play suboptimally—and in a way that varies with their history of play—until the next time they renegotiate. The relevance of bargaining power derives from an assumption of "no-fault disagreement," in which continuation play following disagreement is independent of the manner in which disagreement occurred.

Our theory achieves two goals: (i) providing a coherent account of endogenous agreement and disagreement in repeated games, where bargaining power plays a central role; and (ii) providing a technical apparatus that is straightforward to apply. In Section 2, we illustrate the main ideas in the context of a principal-agent example with imperfect monitoring, similar to that studied by Levin (2003). In the example, the agent is motivated to exert effort on the equilibrium path by variations in continuation utility that are driven by changes in how the parties would behave if they disagreed. The example illustrates how explicitly addressing disagreement behavior clarifies what is attainable under agreement. Notably, we find that equilibrium effort is highest if the agent has all the bargaining power, while no effort can be sustained if the principal has all the bargaining power.

In Section 3 we formally define *contractual equilibrium* as a recursive specification of actions, agreement paths, and disagreement paths satisfying incentive constraints in every action phase and the bargaining solution in every negotiation phase.[1] Section 4 presents our characterization results for two-player games. We prove that if the stage game is finite then a contractual equilibrium always exists, regardless of the discount factor,[2] and that for any given discount factor the welfare level (sum of individual utilities) attained in contractual equilibrium is unique. Since our methods are constructive, they can be used to fully characterize equilibria in applications, via a simple algorithm. We identify necessary and sufficient conditions for efficiency to be attained as the players become sufficiently patient, as well as simple sufficient conditions. We show that the welfare level is maximized when one player or the other has all the bargaining power.

In Section 5 we describe a fully noncooperative model that provides foundations for our hybrid cooperative-noncooperative approach. In this model, actions in the negotiation phase are cheap-talk messages that players use to coordinate on continuation play. Under some intuitive assumptions this model yields the same outcomes as in contractual equilibrium. More precisely, we show that contractual equilibrium can be viewed as a refinement of subgame perfection in this model, where the refinement constrains how cheap-talk statements that the players make during the negotiation phase are classified as either agreements or disagreements, which then translate into coordination on the continuation path.

In Section 6 we extend our analysis to games with more than two players, imperfect public monitoring, and heterogeneous discount factors. Finally, in Section 7 we discuss the potential for further extensions regarding non-stationary environments, costly transfers, imperfect external enforcement, and coalitional bargaining.

1.1 Relation to the literature

A typical infinite-horizon repeated game can have a vast multiplicity of subgame perfect equilibria, and equilibrium payoffs, particularly when players are patient (see Friedman 1971; Fudenberg and Maskin 1986; Fudenberg, Levine, and Maskin 1994). Informally, it seems the players must coordinate very closely in order to select "their" equilibrium from this multitude, so it often makes sense to suppose that they coordinate in a way that

[1] The term "contractual equilibrium" is defined by Watson (2008) in an analogous way for finite games.

[2] This contrasts with "strong renegotiation proofness," for which existence with transferable utility has been guaranteed only if the discount factor is sufficiently high (Baliga and Evans 2000).

benefits all of them. Something like this intuition underlies the common practice, in both theory and applications, of focusing attention on the Pareto frontier of what is attainable in equilibrium. However, payoffs on the Pareto frontier may be supported by the threat of punishments that depart from the Pareto frontier. If players can select continuation equilibria once the game is underway—in the same way that we informally envisioned them doing at the outset—then threats used to support equilibrium may not be credible.

The *renegotiation proofness* literature addresses this problem as follows. Suppose that the players, after each history, may jointly deviate from their planned continuation strategies to another continuation equilibrium in a particular class. A renegotiation-proof equilibrium is a subgame perfect equilibrium away from which the players would never jointly deviate. Different notions of renegotiation proofness arise from different restrictions on the class of equilibria to which players may jointly deviate. The study of renegotiation proofness was initiated by Bernheim and Ray (1989), Farrell and Maskin (1989), and Pearce (1987).[3] Baliga and Evans (2000), Fong and Surti (2009), and Kranz and Ohlendorf (2009) address renegotiation proofness in games with transferable utility.

Renegotiation proofness does not model renegotiation explicitly. Instead, it simply assumes that play in every subgame should be Pareto optimal within some class. In the context of a model with the possibility of disagreement, renegotiation proofness can be conceptualized as requiring play under disagreement to equal that under agreement—so that there is nothing to bargain over. In contrast, we follow the example set by the literature on bargaining, where suboptimal outcomes are allowed off the equilibrium path—and, indeed, are key drivers of equilibrium behavior. Alternating-offer bargaining, for instance, specifies a default outcome should agreement not be reached.[4] Cooperative bargaining concepts, like Nash bargaining, the Shapley value, and the core, take as primitives the payoffs that players or coalitions would receive under disagreement.

In some applications in the literature, there is a tradition of incorporating cooperative negotiation into otherwise noncooperative games. An example is the hold-up-based theory of the firm originated by Grossman and Hart (1986) and Hart and Moore (1990). Much of this literature examines finite-horizon models in which parties form a contract, engage in specific investments, and then have the opportunity to renegotiate their contract before

[3]Renegotiation proofness is further developed by Asheim (1991), Abreu, Pearce, and Stacchetti (1993), Abreu and Pearce (1991), Bergin and MacLeod (1993), and Ray (1994), among others.

[4]In the case of infinite horizon bargaining in which the game continues until agreement, the default option is a path that never terminates the game.

trade takes place.[5] In finite-horizon models, the relationship between cooperative and noncooperative accounts of negotiation is generally straightforward.

Bargaining is also an element in the *relational contracts* literature, which examines how agreements (such as in an employment relationship) can be self-enforced through repeated play. Prominent examples include Radner (1985), MacLeod and Malcomson (1989), Baker, Gibbons, and Murphy (1994, 2002), MacLeod (2003), and Levin (2003). Some models in this literature specify repeated games without conditions on the implications of renegotiation and disagreement, so bargaining power does not play a role. Other models assume that the negotiating parties permanently separate if they disagree, or at least switch permanently to a stage game equilibrium or a spot contract.[6] This approach—assuming that the parties permanently separate following a disagreement—is also commonly used to integrate wage bargaining into macroeconomic models with labor market search frictions, following Diamond and Maskin (1979) and Diamond (1982).[7] These models typically incorporate a cooperative-theory account of negotiation (the generalized Nash solution) and bargaining power determines the outcome relative to the outside options.

In summary, in the existing applied literature bargaining power typically surfaces in the context of a disagreement point that represents discontinuing or irreversibly altering the relationship. Furthermore, the disagreement point is a sunk value that is not influenced by ongoing play. Our model, in contrast, examines ongoing relationships in which the parties can renegotiate in each period and will continue to interact even if they fail to reach an agreement. The disagreement point is endogenously determined by continuation play, and may be history-dependent. Bargaining surplus arises partly because of the players' ability to make transfers at the time of negotiation. Each player's bargaining power derives from the fixed protocol by which negotiation takes place, and from the assumption of no-fault disagreement. Our basic model is a hybrid form in which negotiation is given by a cooperative bargaining solution and stage-game actions are described noncooperatively.

[5]Typically, parties agree on ex-post efficient actions in the renegotiation phase, but if renegotiation fails then the agents obtain the default outcome specified by their initial contract. The key results in this literature describe how the initial contract (which may, for instance, assign ownership of assets) affects the degree to which the parties' investment decisions affect their default outcome, and how the initial contract ought to be structured to encourage efficient investment. The theory is further developed by Hart (1995), Aghion and Bolton (1992), Rajan and Zingales (1998), Hart and Moore (1994), and Aghion and Tirole (1997), among many others.

[6]See, for instance, Baker, Gibbons, and Murphy (1994, 2002) and Schmidt and Schnitzer (1995).

[7]Example are Pissarides (1985, 1987), Hosios (1990), Mortensen and Pissarides (1994), den Haan, Ramey, and Watson (2000), Moscarini (2005), Cahuc, Postel-Vinay, and Robin (2006), and Flinn (2006).

We provide foundations by specifying a fully noncooperative version of the model.[8]

2 Example: A Principal-Agent problem

In this section, we illustrate the basic idea of contractual equilibrium in the context of a principal-agent model with moral hazard, similar to that studied by Levin (2003). We sketch the essential elements of the construction without all of the details. In each period $1, 2, \ldots$, the agent chooses effort $e \in \mathcal{E} = [0, \overline{e}]$, incurring a cost $c(e)$. The principal's profit is a random variable y, with probability density $f(\cdot|e)$ and support $\mathcal{Y} = \left[\underline{y}, \overline{y}\right]$. The principal does not observe e, but y is public. We assume that c is strictly increasing and strictly convex, $c(0) = 0$, f has the monotone likelihood property, and $f\left(y \middle| e = c^{-1}(\cdot)\right)$ is convex. The players can voluntarily pay each other in the negotiation phase,[9] and they share a common discount factor $\delta < 1$. We normalize payoffs by $(1 - \delta)$ to put them in average terms.

Suppose there is no external enforcement, so the principal cannot commit to paying a share of profit. What is the maximum effort that can be sustained in equilibrium, if the principal and the agent engage in Nash bargaining at the start of each period? The answer depends on the allocation of bargaining power, where $\pi_A \geq 0$ is the agent's bargaining share, and $1 - \pi_A$ is the principal's.

The essence of a contractual equilibrium is that in each period the players negotiate an optimal agreement, and they share the surplus relative to what they would attain under disagreement. Under disagreement, there is no transfer and the players may play suboptimally (and differently after different histories). Indeed, their agreement specifies how they should coordinate their future off-path play under disagreement. So to construct a contractual equilibrium we must specify play under disagreement, and then derive play under agreement. In a two-player game such as this, it suffices to construct an equilibrium using a two-state machine, where state 1 rewards the agent and state 2 punishes the agent.

In state 1, for disagreement play we seek an effort level that rewards the agent—e.g., zero effort. Since it is the agent's optimum in the stage game, there is no need for continuation utilities to vary in order to provide him incentives. So after a disagreement in state 1,

[8] Another example of a hybrid cooperative-noncooperative approach with a finite horizon is Brandenburger and Stuart (2007). Also relevant are papers that examine infinite-horizon bargaining with payoff-relevant actions between offers, such as Busch and Wen (1995) and Wolinsky (1987).

[9] Nothing changes if they can also pay each other in the action phase.

the players stay in state 1 the following period regardless of the realized output. Under agreement in state 1, the players recognize that their "outside option" is to disagree, implement zero effort, and then return to state 1 next period. Therefore their utility (specifically, their vector of average discounted utilities) under disagreement is a convex combination of $(0, 0)$ and their agreement utility in state 1. Since the agent obtains a π_A fraction of the surplus, their agreement utility is $(\pi_A, 1 - \pi_A) \cdot (\mathbb{E}(y|e^*) - c(e^*))$, where e^* is the equilibrium path effort. To attain this utility, the principal makes a payment to the agent as part of their agreement. The principal is willing to pay because doing so gives her strictly higher utility than does the disagreement that would arise should she fail to pay.

In state 2, we seek an effort level under disagreement that punishes the agent. The best candidate is the equilibrium path effort e^*, which is the highest effort that can be enforced using equilibrium continuation utilities. This effort level is enforced by staying in state 2 for low realizations of y, and transiting to state 1 for high realizations of y. As in Levin (2003), the optimal cutoff for enforcing any effort level e is $\hat{y}(e)$, the output level at which the likelihood ratio for effort e, as a function of y, crosses zero. In fact, this same incentive scheme must be used whenever the agent is to exert effort e^*, i.e., in state 1 under agreement as well as in state 2 under both agreement and disagreement.

Since optimal effort is played under disagreement in state 2, the disagreement utility is already on the Pareto frontier of what is attainable in equilibrium. Hence when the players negotiate, there is no surplus to share, and so play is the same under agreement and disagreement. Because the agent can always deviate to zero effort in every period, his utility in state 2 must be at least zero. To provide the maximal incentives, in fact his utility in state 2 should be exactly zero.

Equilibrium path effort e^* is thus a fixed point of the agent's optimization problem:

$$e^* \in \arg\max_{e \in \mathcal{E}} \Big[-(1-\delta)\, c(e) + \delta \Pr\big(y \geq \hat{y}(e^*) \big| e\big) \pi_A \big(\mathbb{E}(y|e^*) - c(e^*)\big) \Big] \qquad (1)$$

Because the principal and the agent negotiate over how to play, they will jointly select the highest fixed point. Since incentives are stronger the more weight the agent places on the second term in his objective function, we see that e^* increases in δ and π_A.

Observe that only zero effort is supported if the principal has all the bargaining power ($\pi_A = 0$). Since the principal also receives zero utility if she has no bargaining power ($\pi_A = 1$), her utility is non-monotone in π_A. In this way, contractual equilibrium allows us to investigate the role of bargaining power in Levin's relational contracting environment. The key difference between contractual equilibrium and Levin's "strongly optimal" equilibrium is that the former incorporates a theory of disagreement (including the assumption of no-fault disagreement) that sets a default point for negotiation each period. Also, we allow the agents to endogenously select any division of the surplus (and in equilibrium select the division that satisfies Nash's bargaining solution). Levin's results rely on continuation play that, following an out-of-equilibrium offer by the principal, punishes the principal if the agent rejected but punishes the agent if the agent accepted. Since the disagreement outcome is sensitive to the manner of disagreement, there is little role for the exercise of bargaining power.

Note that if the parties were required to permanently separate after a disagreement, then they would be unable to support positive effort in a contractual equilibrium. Every negotiation along the equilibrium path would lead to the same agreement regardless of the history, giving the agent no incentive to exert effort. The key difference is that under contractual equilibrium the disagreement point may vary with the history, so that the agreement also varies with the history, providing incentives.

On the other hand, if we alter the model so that transfers occur simultaneously with stage-game actions and there is no actual negotiation phase, patient parties can support positive effort in a strongly renegotiation-proof equilibrium (Farrell and Maskin 1989), as shown by Baliga and Evans (2000), although bargaining power would then not play a role. When the principal reneges on a payment in one period, such an equilibrium disciplines her by exacting a large fine in the next period while the agent exerts zero effort. If we assume that transfers are allowed in the negotiation phase as well as the action phase, then such a punishment is not sustainable under contractual equilibrium. In the negotiation phase the players would agree to behave efficiently instead, and use an immediate transfer to split the surplus. The key difference is that most approaches to renegotiation proofness do not allow for transfers during renegotiation, whereas in contractual equilibrium the negotiation-phase transfer "seals the deal."

3 Contractual equilibrium

3.1 Basics

We consider two-player, repeated games with renegotiation and transfers. A game in this class is defined by a stage game $\langle A, u \rangle$, a discount factor δ (by which the players exponentially discount their payoffs across periods), and bargaining weights $\pi = (\pi_1, \pi_2)$. Here $A \equiv A_1 \times A_2$ is the space of action profiles and $u : A \to \mathbb{R}^2$ is the stage-game payoff function. Assume that $\delta \in (0, 1)$, and $\pi \in \mathbb{R}^2$ satisfies $\pi_1 + \pi_2 = 1$ and $\pi_1, \pi_2 \geq 0$. We express repeated game utilities in discounted average terms to facilitate comparison to stage game payoffs.

The players, $i \in \{1, 2\}$, interact over an infinite number of discrete periods. In each period there are two phases of interaction: a *negotiation phase* and an *action phase*. Prior to the negotiation phase, the players can observe an arbitrary public correlation device. In the negotiation phase, the players can either disagree, or jointly agree on their continuation strategies and make immediate, voluntarily monetary transfers. Money enters their payoffs quasi-linearly. In the action phase, the players simultaneously choose actions in the stage game. We analyze behavior in the negotiation phase using the generalized Nash bargaining solution and we study individual incentives in the action phase, so we have a hybrid model.

By distinguishing between the negotiation and action phases, we can articulate theories of disagreement and agreement. The theory of disagreement accounts for how the players coordinate from the action phase of the current period in the event that they fail to reach an agreement today. The theory of agreement describes how negotiation is resolved relative to the disagreement point. We consider a simple disagreement theory in which, if the players fail to reach an agreement then (i) no transfers are made and (ii) the players coordinate on a predetermined action profile for the current period and a specification of continuation values from the following period that together form an equilibrium from the action phase. The disagreement point may be a function of the history of interaction through the previous period.

For simplicity, and without loss of generality, we define the game so that the players negotiate directly over their continuation payoffs. A strategy profile in this game is described by two mappings: one from histories ending in realizations of the public randomization device to the agreement made in the negotiation phase, and one from histories ending with an agreement or a disagreement to mixed action profiles in the stage game. An agreement comprises a monetary transfer and a continuation payoff vector.

A *contractual equilibrium* is a strategy profile in which, after every history, behavior in the negotiation phase is consistent with the Nash bargaining solution, the players select sequentially rational actions in the action phase, and they attain (in expectation) the continuation payoffs that they agree upon. In a Nash bargain, they choose an agreement that maximizes their welfare level (the sum of their utilities), and they use an immediate transfer to divide the surplus of negotiation (the difference in welfare levels between the agreement and the disagreement) according to the bargaining weights π_1 and π_2. That is, player i obtains his disagreement payoff plus a π_i fraction of the surplus. Since our recursive construction defines strategies implicitly and without loss of generality (cf. Abreu, Pearce, and Stacchetti 1990), we refrain from introducing the extra notation necessary to define strategies explicitly.

3.2 Recursive construction

We next define contractual equilibrium recursively, and provide some basic characterizations. We start with some notation. Let $\mathbb{R}^2_0 \equiv \{m \in \mathbb{R}^2 | m_1 + m_2 = 0\}$ be the set of budget-balanced transfer vectors. For any product set H, let $\Delta^{\mathrm{U}} H$ denote the set of probability distributions over H that are uncorrelated across dimensions. Let $V \subset \mathbb{R}^2$ denote a set of continuation values (vectors) from the start of a given period, prior to the realization of the public randomization device. The dynamic programming equation characterizing an agreement is

$$y = (1-\delta)\big(m + u(\alpha)\big) + \delta g(\alpha), \tag{2}$$

where y is the continuation value that is agreed upon, m is a budget-balanced transfer, $\alpha \in \Delta^{\mathrm{U}} A$ is the mixed action profile to be played in the current period, and $g : A \to V$ gives the continuation value from the start of the next period as a function of current-period actions in the stage game. Here both u and g are extended to the space of mixed actions $\Delta^{\mathrm{U}} A$ by taking expectations. Under disagreement, there are no immediate transfers. So the dynamic program for a disagreement is

$$\underline{y} = (1-\delta)u(\alpha) + \delta g(\alpha). \tag{3}$$

In equilibrium, the players' actions in any given period must be sequentially rational, whether under agreement or disagreement. That is, they must play a Nash equilibrium in the game implied by the dynamic program.

Definition 1. The function $g : A \to V$, extended to $\Delta^{\mathrm{U}} A$, *enforces* $\alpha \in \Delta^{\mathrm{U}} A$ if α is a Nash equilibrium of $\langle A, (1-\delta)u + \delta g \rangle$.

For an arbitrary set V of continuation values available from the following period, the set of payoffs that can be supported in the current negotiation phase is given by the operator C:

$$C(V) \equiv \left\{ y \in \mathbb{R}^2 \middle| \exists\ m \in \mathbb{R}_0^2,\ g : A \to V, \text{ and } \alpha \in \Delta^{\mathrm{U}} A \text{ s.t. } g \text{ enforces } \alpha \text{ and Eq. 2 holds} \right\}. \tag{4}$$

Note that $C(V)$ is closed under transfers, meaning that $y \in C(V)$ and $m \in \mathbb{R}_0^2$ imply that $y + m \in C(V)$. Also, if V is compact then it must be that $\max\{y_1 + y_2 | y \in C(V)\}$ exists, due to the graph of the Nash equilibrium correspondence being closed.

Under disagreement, there are no transfers, so the set of payoffs that can arise following a disagreement is given by the operator D:

$$D(V) \equiv \left\{ \underline{y} \in \mathbb{R}^2 \middle| \exists\ g : A \to V \text{ and } \alpha \in \Delta^{\mathrm{U}} A \text{ s.t. } g \text{ enforces } \alpha \text{ and Eq. 3 holds} \right\}. \tag{5}$$

Since the Nash equilibrium graph is closed, D preserves compactness.

The outcome of negotiation must satisfy the generalized Nash bargaining solution which, for any particular $\underline{y}$, implies maximizing the players' joint value (the sum of their continuation payoffs) and dividing the surplus according to their bargaining weights. That is, the solution y satisfies

$$y_1 + y_2 = \max_{y' \in C(V)} y_1' + y_2' \text{ and } \pi_2(y_1 - \underline{y}_1) = \pi_1(y_2 - \underline{y}_2). \tag{6}$$

The second part of this condition is equivalent to

$$y_i = \underline{y}_i + \pi_i(y_1 + y_2 - \underline{y}_1 - \underline{y}_2) \text{ for each player } i. \tag{7}$$

Given a set of possible agreements Y and disagreements $\underline{Y}$, the set of continuation values that can arise following realization of the public randomization device is given by the operator B:

$$B(Y, \underline{Y}) \equiv \left\{ y \in Y \middle| \exists\ \underline{y} \in \underline{Y} \text{ s.t. Eq. 6 holds} \right\}. \tag{8}$$

Note that B is well defined as long as $\max_{y \in Y}\{y_1 + y_2\}$ exists and Y is closed under constant transfers. If $\underline{Y}$ is compact, then $B(Y, \underline{Y})$ is also compact. Observe that if $\underline{Y}$ is nontrivial then $B(Y, \underline{Y})$ is multivalued. That is, the model has fixed bargaining weights, but coordination on different disagreement points gives rise to different bargaining outcomes.

We next put the three conditions together to describe the formal relation between V, $C(V)$, and $D(V)$. Because the repeated game is stationary, we require that these sets apply to all periods. Consequently, if the sets describe equilibria with negotiation in each period, it must be that $V = \text{co}\, B(C(V), D(V))$, where "co" denotes "convex hull" to account for the public randomization device. We utilize the standard recursive formulation of Abreu, Pearce, and Stacchetti (1990) to describe sets that satisfy the required properties.

Definition 2. A set $V \subset \mathbb{R}^2$ satisfies *negotiated self-generation* **(NSG)** if $V = \text{co}\, B(C(V), D(V))$.

When discussing an NSG set V, we will refer to $Y = C(V)$ and $\underline{Y} = D(V)$ without reference to C or D. The next lemma follows from the properties of B, C, and D already discussed.

Lemma 1. *If $V \subset \mathbb{R}^2$ satisfies NSG then it is a convex line segment with slope -1 and finite length. Thus, all points in V have the same joint value for the players.*

Thus, any NSG set is defined by its endpoints. For most of the presentation, we focus on the case in which V contains its endpoints and is therefore closed; the case in which V does not contain its endpoints is handled in the proofs. We denote the upper-left endpoint (giving the lowest payoff for player 1) z^1 and we will let z^2 be the endpoint favoring player 1. Clearly, z^1 and z^2 are extreme points of $B(Y, \underline{Y})$. Therefore, there are disagreement points $\underline{y}^1$ and $\underline{y}^2$ relative to which z^1 and z^2 are the bargaining outcomes satisfying Eq. 6.

Definition 3. A set $V \subset \mathbb{R}^2$ **dominates** another set $\hat{V} \subset \mathbb{R}^2$ if for every $\hat{v} \in \hat{V}$ there exists some $v \in V$ such that $v_1 \geq \hat{v}_1$ and $v_2 \geq \hat{v}_2$.

Note that $\hat{V} \subset V$ implies that V dominates $\hat{V}$. In terms of the endpoints, if (z^1, z^2) characterizes V and $(\hat{z}^1, \hat{z}^2)$ characterizes $\hat{V}$, then V dominates $\hat{V}$ if and only if $z_2^1 \geq \hat{z}_2^1$ and $z_1^2 \geq \hat{z}_1^2$. Our main definition combines negotiated self-generation with dominance to form a notion of equilibrium with negotiation.

Definition 4. A set $V^* \subset \mathbb{R}^2$ is a *contractual equilibrium value (CEV) set* if it satisfies NSG and it dominates every other NSG set.

Note that by definition if two NSG sets dominate each other, they must be the same set. Hence there can be at most one set that represents contractual equilibrium. However, there can be many strategy profiles consistent with this set. In our hybrid cooperative-noncooperative model, a contractual equilibrium strategy profile can be expressed as a selection $(y, \underline{y}) \in C(V^*) \times D(V^*)$ for each $v \in V^*$, a selection $\alpha \in \Delta^{\text{U}} A$ for each $y \in C(V^*) \cup D(V^*)$, and a selection $m \in \mathbb{R}_0^2$ for each $\underline{y} \in D(V^*)$ that are consistent with the forgoing construction. The selections of y, $\underline{y}$, and m represent joint actions, while the selection of α represents a vector of individual actions.

4 Results

4.1 Existence and characterization

Our main result establishes existence, and thus uniqueness, of the CEV set for any discount factor.

Theorem 1. *Consider any two-player repeated game with renegotiation and transfers, defined by* $\langle A, u, \delta, \pi \rangle$*. If* A *is finite and* $\delta \in [0, 1)$ *then the game has a unique CEV set* V^**.*

The rest of this section contains the proof of the theorem, which also shows how to construct the set V^*. For most of the analysis we constrain attention to NSG sets that are closed—in other words, lines that contain their endpoints. We argue at the end that open sets are dominated.

The first step in the analysis is a characterization of the endpoints of an NSG set. Consider a closed NSG set V, so it contains its endpoints z^1 and z^2. We will express z^1 and z^2 in relation to optimization problems parameterized by $z^2_1 + z^2_2$ and $z^2_1 - z^1_1$. Note that $z^2_1 + z^2_2$ is the joint value of the relationship (in average terms) and is equal to $z^1_1 + z^1_2$. We call $z^2_1 - z^1_1$ the *payoff span*. We will show the steps for the characterization of z^2; the logic is the same for z^1.

Since V is an NSG set, and z^2 is the endpoint that most favors player 1, z^2_1 is the maximum payoff for player 1 that can be supported utilizing continuation values from V associated with the next period:

$$z^2_1 = \max \left\{ v_1 \middle| v \in B\big(C(V), D(V)\big) \right\} \tag{9}$$

Because elements in $B(C(V), D(V))$ correspond to various disagreement points in $D(V)$, this maximization problem can be expressed using the bargaining solution (Eq. 6), taking the disagreement point as the choice variable. Letting

$$L \equiv \max_{y \in C(V)} y_1 + y_2, \tag{10}$$

note that, for any given disagreement point $\underline{y} \in D(V)$, the Nash bargaining solution gives player 1 the value $= \underline{y}_1 + \pi_1(L - \underline{y}_1 - \underline{y}_2)$, which equals $\pi_2 \underline{y}_1 - \pi_1 \underline{y}_2 + \pi_1 L$. Eq. 9 can therefore be written as:

$$
\begin{aligned}
z_1^2 = \max_{\underline{y},g,\alpha} \;\; & \pi_2 \underline{y}_1 - \pi_1 \underline{y}_2 + \pi_1 L, \\
\text{s.t.} \;\; & \begin{cases} \underline{y} = (1-\delta)u(\alpha) + \delta g(\alpha), \\ g : A \to V \text{ enforces } \alpha. \end{cases}
\end{aligned}
\tag{11}
$$

We next rewrite the optimization problem with a change of variables. Define $\eta(a) \equiv g_1(a) - z_1^2$ for every a. Because the welfare level of every point in V is $z_1^2 + z_2^2$, we have $g_2(a) = z_2^2 - \eta(a)$. Also, the constraint that $g(a) \in \text{co}\{z^1, z^2\}$ is equivalent to the requirement that $\eta(a) \in [z_1^1 - z_1^2, 0]$. Using η to substitute for g and $\underline{y}$ and combining terms, we see that Eq. 11 is equivalent to:

$$
\begin{aligned}
z_1^2 = \max_{\eta,\alpha} \;\; & (1-\delta)\big[\pi_2 u_1(\alpha) - \pi_1 u_2(\alpha)\big] + \delta\big(\pi_2 z_1^2 - \pi_1 z_2^2\big) + \delta\eta(\alpha) + \pi_1 L, \\
\text{s.t.} \;\; & \begin{cases} \eta : A \to [z_1^1 - z_1^2, 0] \text{ extended to } \Delta^{\mathrm{U}} A, \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, (1-\delta)u + \delta(\eta, -\eta) + \delta z^2 \big\rangle. \end{cases}
\end{aligned}
\tag{12}
$$

Using this equality we solve for z_1^2 in terms of $z_1^2 + z_2^2$ and $z_1^2 - z_1^1$ (the joint value and span of V). The calculations produce the following implicit characterization of z_1^2:

$$
z_1^2 = \pi_1\big[z_1^2 + z_2^2\big] + \overline{\gamma}(z_1^2 - z_1^1), \tag{13}
$$

where

$$
\begin{aligned}
\overline{\gamma}(d) \equiv \max_{\eta,\alpha} \;\; & \pi_2 u_1(\alpha) - \pi_1 u_2(\alpha) + \frac{\delta}{1-\delta}\eta(\alpha), \\
\text{s.t.} \;\; & \begin{cases} \eta : A \to [-d, 0], \text{ extended to } \Delta^{\mathrm{U}} A, \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, (1-\delta)u + \delta(\eta, -\eta) \big\rangle. \end{cases}
\end{aligned}
\tag{14}
$$

By similar calculations, we obtain an implicit characterization of z_1^1:

$$
z_1^1 = \pi_1\big[z_1^2 + z_2^2\big] + \underline{\gamma}(z_1^2 - z_1^1), \tag{15}
$$

where

$$\underline{\gamma}(d) \equiv \min_{\eta,\alpha} \ \pi_2 u_1(\alpha) - \pi_1 u_2(\alpha) + \frac{\delta}{1-\delta}\eta(\alpha),$$
$$\text{s.t.} \begin{cases} \eta : A \to [0, d], \text{ extended to } \Delta^{\mathrm{U}} A, \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, (1-\delta)u + \delta(\eta, -\eta)\big\rangle. \end{cases} \tag{16}$$

We know that the optima defining $\overline{\gamma}$ and $\underline{\gamma}$ exist because the stage game is finite, the set of feasible η functions is compact, and the Nash correspondence is upper hemi-continuous.

Now we can compare NSG sets by using the functions $\underline{\gamma}$ and $\overline{\gamma}$. We find that the NSG sets are ranked by dominance.

Lemma 2. *Suppose that V and $\hat{V}$ are both NSG and closed. Let z^1 and z^2 be the endpoints of V. Let $\hat{z}^1$ and $\hat{z}^2$ be the endpoints of $\hat{V}$. If $z_1^2 - z_1^1 \geq \hat{z}_1^2 - \hat{z}_1^1$ then V dominates $\hat{V}$.*

Proof. Suppose that $z_1^2 - z_1^1 \geq \hat{z}_1^2 - \hat{z}_1^1$. The larger payoff span of V can support weakly more mixed actions in the stage game as equilibria than can $\hat{V}$ simply because V allows for a greater range of continuation values in the next period. This comparison does not depend on the location of the endpoints or the joint values of the two sets (which amount only to constants in the players' payoffs), only their relative payoff spans. Thus, any mixed action that can be supported in the context of $\hat{V}$ can also be supported in the context of V. This implies that $z_1^2 + z_2^2 \geq \hat{z}_1^2 + \hat{z}_2^2$; that is, the welfare level of V weakly exceeds the welfare level of $\hat{V}$. Furthermore, since $\overline{\gamma}(d)$ is increasing in d, Eq. 13 and the larger span of V imply that $z_1^2 \geq \hat{z}_1^2$. Also note that, using $\pi_1 + \pi_2 = 1$ and $z_1^1 + z_2^1 = z_1^2 + z_2^2$, we can rearrange Eq. 15 to form $z_2^1 = \pi_2(z_1^2 + z_2^2) - \underline{\gamma}(z_1^2 - z_1^1)$. Since $\underline{\gamma}$ is decreasing, we conclude that $z_2^1 \geq \hat{z}_2^1$, which is sufficient to establish that V dominates $\hat{V}$. □

We next use the functions $\overline{\gamma}$ and $\underline{\gamma}$ to prove the existence of a dominant NSG set. The difference

$$\Gamma(d) \equiv \overline{\gamma}(d) - \underline{\gamma}(d) \tag{17}$$

will be of particular interest. Note that Γ maps the payoff span of continuation values from the next period into the supported payoff span from the beginning of the current period.

Observe that every NSG set $V = \mathrm{co}\{z^1, z^2\}$ is associated with a fixed point of Γ in that, for the payoff span $d = z_1^2 - z_1^1$, we have $d = \Gamma(d)$. We need to show that Γ has a maximal fixed point d^* and then construct V^* from it. To this end, observe that Γ is increasing because larger payoff spans relax the constraints in the problems that define $\overline{\gamma}$ and $\underline{\gamma}$. It is also bounded because u is bounded and δ is fixed. By Tarski's fixed-point theorem, we therefore know that Γ has a maximal fixed point d^*. To find the associated NSG set V^*, we simply calculate:

$$L^* \equiv \max_{\eta,\alpha} \; u_1(\alpha) + u_2(\alpha), \qquad \text{s.t.} \begin{cases} \eta : A \to [0, d^*], \text{ extended to } \Delta^{\mathrm{U}} A, \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, (1-\delta)u + \delta(\eta, -\eta) \big\rangle. \end{cases} \tag{18}$$

Then we obtain z_1^{2*} and z_1^{1*} using Eq. 13 and Eq. 15, with L^* in place of $z_1^2 + z_2^2$ and d^* in place of $z_1^2 - z_1^1$. Finally, we have $z_2^{2*} \equiv L^* - z_1^{2*}$ and $z_2^{1*} \equiv L^* - z_1^1$. We have thus have identified points z^{1*} and z^{2*} and can define $V^* \equiv \mathrm{co}\{z^{1*}, z^{2*}\}$. By construction, V^* is NSG and it dominates all other NSG sets.

We finish the proof by addressing the case of an open set, where V does not contain one or both endpoints. Taking the closure does not necessarily form an NSG set, because new Nash equilibria could emerge in the game implied by the dynamic program. However, V^* is closed and we can see that, using the arguments already employed, the joint value of V must be weakly lower than the that of V^*. In the case in which the joint values are the same, we have $V \subset V^*$ and so V^* dominates. Otherwise, the comparison of endpoints yields the dominance relation.

4.2 Efficiency

Inspection of $\overline{\gamma}$ and $\underline{\gamma}$ reveals that Γ is bounded and increasing in δ. In fact, we can show that if $\Gamma(d) = \varepsilon > 0$ for some d and some δ, then $\Gamma(\varepsilon) \geq \varepsilon$ for large enough discount factors, which means that d^* is bounded away from zero for δ close enough to one. This implies that any action profile can be supported in a single period when players are patient, and we have the following result.

Theorem 2. *For a given repeated game with renegotiation and transfers, if* $\overline{\gamma}(\infty) > \underline{\gamma}(\infty)$ *then* $d^* > 0$ *and* V^* *is a subset of the efficient frontier for* δ *sufficiently close to 1. If* $\overline{\gamma}(\infty) = \underline{\gamma}(\infty)$ *then* V^* *attains the same welfare level as the welfare-maximizing Nash equilibrium of the stage game.*

Example 1 (The Prisoners' Dilemma). The result is illustrated by the Prisoners' Dilemma (Figure 1). The stage game and feasible payoff set are pictured in Figure 1. To apply

1 \ 2	C	D
C	1, 1	-r, x
D	x, -r	0, 0

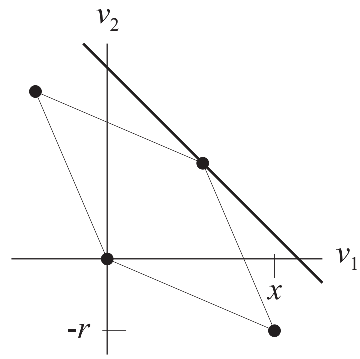

FIGURE 1. THE PRISONERS' DILEMMA WITH TRANSFERS. The stage game parameters satisfy $x > r > 0$. The Pareto frontier is a line with slope -1 through the point $(1, 1)$.

Theorem 2, we observe that

$$\overline{\gamma}(\infty) = \max\{\pi_2 - \pi_1 - (x-1), \pi_2 x + \pi_1 r - r, -\pi_2 r - \pi_1 x, 0\}, \tag{19}$$

$$\underline{\gamma}(\infty) = \min\{\pi_2 - \pi_1 + (x-1), \pi_2 x + \pi_1 r, -\pi_2 r - \pi_1 x + r, 0\}, \tag{20}$$

where the elements in each maximand are the values at CC, DC, CD, and DD, in that order. Because D is a dominant strategy in the stage game, the values at mixed strategies are convex combinations of these. It is clear that $\underline{\gamma}(\infty) > \overline{\gamma}(\infty)$ if and only if $x > r$, which by Theorem 2 occurs if and only if CC is supported in contractual equilibrium for sufficiently high $\delta < 1$. If, on the other hand, $x \leq r$, then $V^* = \{(0,0)\}$ regardless of δ.

Suppose that in fact $x > r$. To get a feel for the construction of the set V^*, consider player 1's favorite point z^2. It will be the case that z^2 is achieved with reference to the disagreement point $\underline{y}^2$ that is furthest in the direction $(\pi_2, -\pi_1)$. This disagreement point, in turn, will be the weighted average of $u(\text{DC}) = (x, -r)$ and a point $v' \in V^*$. The former gets the weight $(1-\delta)$ and represents the payoff in the current period, whereas the latter gives the continuation payoff from the next period and has the weight δ. It is clearly best to "push" v' down and to the right, to favor player 1. However, v' cannot equal z^2 because we need room to punish player 2 if he were to deviate from DC. Thus, we specify that the players select DC in the current period. If player 2 cheats, the continuation value z^2 is selected; otherwise, the continuation value is v'. The construction is pictured in Figure 2. Note that $\underline{y}^2$ is the disagreement point and it leads to z^2 as the solution to the bargaining problem.

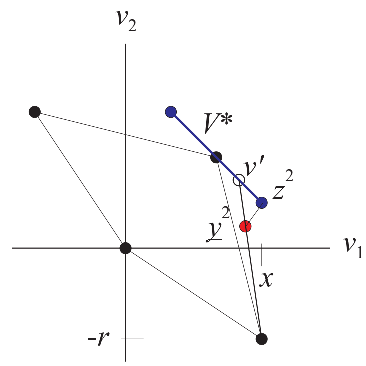

FIGURE 2. PRISONERS' DILEMMA z^2 CONSTRUCTION. The endpoint z^2 is attained by playing CC in the stage game and using a transfer to split the surplus relative to the disagreement point $\underline{y}^2$. The disagreement point $\underline{y}^2$, in turn, is attained by playing DC in the stage game and continuing with promised utility v' if no deviation occurs, and promised utility z^2 if player 2 deviates.

Next we provide methods for constructing contractual equilibria explicitly for a variety of common settings.

Theorem 3. *For each i, let a^i be a pure action profile in the stage game. If a^i_i is a best response to a^i_{-i} for both i, and $(\pi_2, -\pi_1) \cdot \big(u(a^2) - u(a^1)\big) > 0$, then the following two-state automaton, with states $\{\omega^1, \omega^2\}$, yields an NSG payoff set for δ sufficiently high:*

- *Agreement: Play $\arg\max_a \sum_i u_i(a)$. If nobody deviates or both deviate, randomize between the two states with equal probabilities. If player i deviates unilaterally, go to state ω^i.*
- *Disagreement: In state ω^i, play α^i. If nobody deviates or both deviate, stay in state ω^i. If player j deviates unilaterally, go to state ω^j.*

This works because player i needs no incentives at a^i, while player $-i$ can be given strong incentives by the threat of switching to ω^{-i}. So under disagreement, if nobody deviates, they can stay in the same state. This being the case, their agreement utility in each state is on a $\frac{\pi_2}{\pi_1}$-sloped line from their stage game payoff under disagreement. This is illustrated in the following example.

Example 2 (The Prisoners' Battle of the Sexes)*.* In the Prisoners' Battle of the Sexes, illustrated in Figure 3, CC is the optimal stage game profile, but C is dominated by both D and E for each player. There are three stage game Nash equilibria: DE, ED, and a mixed strategy equilibrium. The two-state automaton described in Theorem 3 is implemented by

			C	D	E
				2	
		C	4, 4	0, 5	0, 5
1		D	5, 0	1, 1	2, 4
		E	5, 0	4, 2	1, 1

Figure 3. The Prisoners' Battle of the Sexes game.

selecting $\alpha^1 = \text{DE}$ and $\alpha^2 = \text{ED}$. Then, in each state, $\underline{y}^i$ is a convex combination of $u(\alpha^i)$ and z^i.

Since a player who is being minimaxed is always playing a best response in the stage game, this theorem implies a minimax separation condition that may be easy to check in many games.

Corollary 1. *Let m^i be the pure action minimax payoff profile for player i in the stage game. Suppose that $(\pi_2, -\pi_1) \cdot \left(m^2 - m^1\right) > 0$. Then there exists an efficient contractual equilibrium if the players are sufficiently patient.*

Our next result demonstrates that when the stage game has an interior Nash equilibrium around which the best response functions are differentiable, and an increase in one player's action strictly reduces the other player's stage game payoff, there exists an efficient contractual equilibrium if the players are sufficiently patient.

Theorem 4. *Suppose that $A_i \supset \left[\underline{\alpha}_i, \overline{\alpha}_i\right]$ for both i, and there exists a Nash equilibrium α^* in the interior of $\left[\underline{\alpha}_1, \overline{\alpha}_1\right] \times \left[\underline{\alpha}_2, \overline{\alpha}_2\right]$. Suppose that there also exist $\eta > 0$ and $k < \infty$ such that, for all α in an η-neighborhood of α^* and for all i, $-k < dBR_i(\alpha_{-i})/ds_{-i} < k$ (where BR_i is player i's best response function) and $du_i(\alpha)/d\alpha_{-i} < -1/k$.[10] Then there exists an efficient contractual equilibrium if the players are sufficiently patient.*

As an example, the symmetric Cournot duopoly game does not satisfy the conditions of Corollary 1, since both firms earn zero profits whenever one firm is minimaxed. However, it has an interior Nash equilibrium, its best response functions have bounded slope, and each firm's payoff is decreasing in the other's quantity. That is, it satisfies the conditions of Theorem 4, and therefore has an efficient contractual equilibrium if the firms are sufficiently patient.

[10]If instead $du_i(\alpha)/d\alpha_{-i} > 1/k$, one could simply relabel A_i to satisfy these suppositions without loss of generality.

		2		
		C	D	E
	C	$1, 1$	$-3, 2$	$-3, 2$
1	D	$2, -3$	$.1, .1$	$0, 0$
	E	$2, -3$	$0, 0$	$0, 0$

FIGURE 4. THE PRISONERS' DOUBLE DILEMMA GAME.

4.3 The role of relative bargaining power

In static bargaining games with transferable utility, the allocation of bargaining power typically has no effect on the level of welfare. In contractual equilibrium, however, the bargaining weights play a critical role in determining the payoff span of the CEV set, and hence the welfare level it attains. In particular, the highest welfare level is attained when bargaining power is extremely unequal.

Theorem 5. *In any game, the welfare level is maximized at either $\pi = (0, 1)$ or $\pi = (1, 0)$.*

Intuitively, V^* is the projection of the disagreement points onto the Pareto frontier in the direction of the bargaining shares. Since equal bargaining shares form a vector perpendicular to the Pareto frontier, they minimize the projected distance between the disagreement points. We illustrate this idea with the following example.

Example 3 (The Prisoners' Double Dilemma). An instructive example is the Prisoners' Double Dilemma shown in Figure 4. This game is similar to the Prisoners' Dilemma, and since $x - r < 0$ for similar reasoning it is not possible to support play of CD, CE, DC, or EC under disagreement. When $\pi = \left(\frac{1}{2}, \frac{1}{2}\right)$, playing the two stage game equilibria—DD and EE—under disagreement can support only an NSG set with a payoff span of 0, and thus efficiency is not attainable. However, whenever $\pi \neq \left(\frac{1}{2}, \frac{1}{2}\right)$ it is easy to see that an NSG set supported by playing DD and EE under disagreement has a strictly positive payoff span, and thus efficiency is attainable for sufficiently high δ.

5 Noncooperative foundations

In this section, we provide noncooperative foundations for contractual equilibrium. The main result of this section, Theorem 6, demonstrates that the set of contractual equilibrium payoffs is identical to the set of refined subgame perfect equilibrium payoffs in a fully non-cooperative repeated game with random-proposer, ultimatum-offer, cheap-talk bargaining in each period, where the refinements endow the cheap-talk bargaining statements with endogenous meaning. We start by specifying the noncooperative model. Suppose that each period comprises four phases: (1) the public randomization phase, (2) the bargaining phase, (3) the voluntary transfer phase, and (4) the action phase.

In the action phase, the players play the stage game $\langle A, u \rangle$ as defined in Section 3. In the transfer phase, the players simultaneously make voluntary, non-negative monetary transfers (that is, each player decides how much money to give to the other). The net transfer is denoted $m \in \mathbb{R}^2$. In the bargaining phase, first nature randomly selects one of the players to make a verbal statement. Nature selects player 1 with probability π_1 and selects player 2 with probability π_2. Let k denote the selected player, who is called the "offerer." The offerer selects a statement from some language set Λ. The other player (the "responder") then says "yes" or "no."[11]

Notice that we have replaced the cooperative bargaining solution used earlier with a noncooperative specification. Further, we have distinguished between verbal communication and monetary transfers by having them in separate phases. Actions in the negotiation phase are payoff-irrelevant cheap talk. We shall build a refinement of the set of subgame-perfect equilibria by imposing conditions on how continuation play relates to this communication in the negotiation phase. That is, we assume that the communication has some intrinsic meaning.

We assume that the language Λ is large enough so that each player can use it to suggest to the other how to coordinate their play in the continuation of the game. One way of ensuring that the language is large enough would be to assume that Λ contains descriptions of all continuation strategy profiles in the game. Since such a construction would be circular (strategies would specify history-dependent statements in the negotiation phase, and these statements would include the description of a strategy profile), it would lead to an interesting technical issue regarding whether an appropriate "universal language" exists.

A simpler approach is to assume that Λ contains the space of possible continuation payoff vectors from the action phase, $\mathbb{R}^2$, and that a selection from this space is viewed by

[11] There is nothing special about the random-proposer ultimatum protocol. Similar results would arise under other bargaining protocols. In particular, the same results arise under standard random-proposer protocols in "stop time" (where all rounds of offers and responses happen in infinitesimal time compared to the length of a period) with exogenous risk of breakdown. See Binmore, Rubinstein, and Wolinsky (1986).

the players as a suggested continuation value. Thus, we define:

$$\Lambda \equiv \{\text{“discontinue”}\} \cup (\mathbb{R}^2 \times \mathbb{R}_0^2). \tag{21}$$

So far, statements in Λ have no particular meaning, but we will make assumptions below that imply meaning in equilibrium. Here is the flavor of what follows: The statement “discontinue” expresses that the offerer will not continue with the current agreement path. A statement of some point $\lambda = (w, m) \in \mathbb{R}^2 \times \mathbb{R}_0^2$ expresses a specific suggestion about how the players should coordinate in the continuation of the game. By stating $\lambda = (w, m)$, the offerer is suggesting that a voluntary monetary transfer of m be made and then the players select continuation strategies that achieve the continuation value w from the action phase in the current period. If the suggestion is feasible and followed, the continuation value would be $y = (1 - \delta)m + w$ from the negotiation phase.

Let S be the set of all subgame-perfect equilibria of our fully noncooperative model, in behavioral strategies. We shall refine S through a series of steps that add meaning to the language. First, we limit attention to the subset S^{C} that satisfies *meaningful agreement*. Meaningful agreement requires that, for every history to the negotiation phase of a given period, the equilibrium continuation satisfies the following: The offerer makes a statement $\lambda = (w, m) \in \mathbb{R}^2 \times \mathbb{R}_0^2$, the responder says “yes,” voluntary transfer m is actually made, and the continuation from the action phase yields the value w. We thus want to think of statements of (w, m) and “yes” as indicating that the players have accepted the equilibrium proposal (w, m) and coordinate to implement it. Because the actions in the bargaining phase are cheap talk and there is a public randomization device (so the players do not need to use communication to jointly randomize), the payoffs supported by S and S^{C} are the same.

The next steps involve building meaningful notions of disagreement and agreement when players deviate from their equilibrium statements. This is done in stages, successively refining the set of equilibria. We describe the events by referring to the set of histories through the transfer phase of a given period, denoted by $\hat{H}$ and defined as follows. Letting Ω denote the set of outcomes of the public randomization device, for any $t \geq 1$ a full t-period history is an element of

$$H^t \equiv (\Omega \times \{1, 2\} \times \Lambda \times \{\text{“yes”}, \text{“no”}\} \times \mathbb{R}_0^2 \times A)^t. \tag{22}$$

Here $\{1,2\}$ refers to the outcome of Nature's random selection of which player is the offerer. Let H^0 denote the null history at the start of the game, and let $H \equiv \cup_{t=0}^{\infty} H^t$. Then

$$\hat{H} \equiv H \times \Omega \times \{1,2\} \times \Lambda \times \{\text{“yes”}, \text{“no”}\} \times \mathbb{R}_0^2. \tag{23}$$

For $\hat{h} \in \hat{H}$ we write $\hat{h} = (h; \omega, k, \lambda, r, m)$, which is $h \in H$ appended with ω (public random draw), k (identity of the offerer), λ (offerer's statement), r (responder's statement), and m (realized voluntary transfer).

Classify a history $\hat{h} \in \hat{H}$ as *ending with a disagreement event* if any of the following occurred at the last period represented by $\hat{h}$: if the offerer says "discontinue," the responder says "no," or the transfer actually made is different from the one named by the offerer. Let $\hat{H}^{\mathrm{D}}$ denote the subset of $\hat{H}$ which end with a disagreement event.

We restrict attention to strategies in which play following a disagreement event does not depend on the manner in which disagreement occurred. We say that two histories $\hat{h}, \hat{h}' \in \hat{H}$ *agree up to the current offerer selection* if $\hat{h} = (h; \omega, k, \lambda, r, m)$ and $\hat{h}' = (h; \omega, k', \lambda', r', m')$ for some common $h \in H$ and $\omega \in \Omega$.

Definition 5 (No-fault disagreement)**.** A strategy profile $s \in S^{\mathrm{C}}$ satisfies *no-fault disagreement* if, for every pair of histories $\hat{h}, \hat{h}' \in \hat{H}^{\mathrm{D}}$ that agree up to the current offerer selection, the continuation specified by s is the same following $\hat{h}$ and $\hat{h}'$. The associated continuation value from the action phase is the *disagreement point.*

Let S^{D} be the subset of S^{C} that satisfies no-fault disagreement.

We next classify some histories as ending in an agreement to adopt a continuation that is supported by playing as if switching to some other history (under the same strategy profile). Formally, for any strategy profile s we define the set $\hat{H}^{\mathrm{I}}(s) \subset \hat{H}$ as follows. Consider any history $\hat{h} = (h; \omega, k, \lambda, r, m) \in \hat{H}$. Then $\hat{h}$ is included in $\hat{H}^{\mathrm{I}}(s)$ if and only if:

1. $\lambda = (w, m') \in \mathbb{R}^2 \times \mathbb{R}_0^2$,

2. $r =$ "yes",

3. $m = m'$ (so that the suggested transfer was made), and

4. there is a "matching history" $\hat{h}' \in \hat{H}$ such that, with the strategy profile s, the continuation value from history $\hat{h}'$ is exactly w.

In other words, an agreement history arises if the offerer suggests a continuation value that is consistent with some other history under the players' strategy profile, if the responder says "yes," and if the suggested transfer is actually made.

We impose the following agreement condition, which says that if the players agree to switch to another continuation that is consistent with their strategy profile, then this actually occurs.

Definition 6 (Agreement-internal)**.** A strategy profile $s \in S^{\mathrm{D}}$ satisfies the **agreement-internal** condition if for every $\hat{h} \in \hat{H}^{\mathrm{I}}(s)$ there is a matching history $\hat{h}' \in \hat{H}$ such that the continuation strategies of s conditional on $\hat{h}$ and $\hat{h}'$ are identical.

Let S^{I} be the subset of S^{D} that satisfy the agreement-internal condition. If an equilibrium is in S^{I}, then the players can never, in any bargaining phase, strictly prefer to "restart" their relationship at a different history—since if such an agreement were suggested, it would indeed be adopted.

We next classify some histories as ending in an agreement to adopt a continuation that is supported by playing as if switching to some other history under a different strategy profile in S^{I}. For a given set $S' \subset S$, let $\bar{L}$ be the supremum of levels (joint continuation values) over all strategies in S' and all histories in H. That is, letting $\bar{v}(s,h)$ denote the continuation value under strategy s following history h,

$$\bar{L} \equiv \sup\{\bar{v}_1(s,h) + \bar{v}_2(s,h) \mid s \in S',\, h \in H\}. \tag{24}$$

We say that some $s^* \in S'$ is a *ranking strategy profile* if, for every $h \in H$, $\bar{v}_1(s^*,h) + \bar{v}_2(s^*,h) = \bar{L}$. Clearly, a ranking strategy profile may not exist for an arbitrary S'.

For any strategy profile s and a set $S' \subset S$ we define the set $\hat{H}^{\mathrm{E}}(s,S') \subset \hat{H}$ as follows. Consider any history $\hat{h} = (h;\omega,k,\lambda,r,m) \in \hat{H}$. Then $\hat{h}$ is included in $\hat{H}^{\mathrm{E}}(s,S')$ if and only if:

1. $\lambda = (w, m') \in \mathbb{R}^2 \times \mathbb{R}^2_0$,
2. $r =$ "yes",
3. $m = m'$, and
4. there is a "matching strategy" $s' \in S'$ and a "matching history" $\hat{h}' \in \hat{H}$, such that the continuation value from history $\hat{h}'$ under strategy profile s' is exactly w, and s' is a ranking strategy profile.

In other words, an agreement history (relative to S') arises if the offerer suggests a continuation value that is consistent with a ranking strategy profile in S', if the responder says "yes," and if the suggested transfer is actually made. Clearly $\hat{H}^{\mathrm{E}}(s, S')$ is limited by the requirement that the matching strategy be ranking, and thus $\hat{H}^{\mathrm{E}}(s, S')$ may be empty.

The second agreement condition says that if the players agree to switch to another continuation that is consistent with a ranking strategy profile in S^{I}, then this actually occurs.

Definition 7 (Agreement-external)**.** A strategy profile $s \in S^{\mathrm{I}}$ satisfies the **agreement-external** condition if for every $\hat{h} \in \hat{H}^{\mathrm{E}}(s, S^{\mathrm{I}})$ there is a matching strategy s' and a matching history $\hat{h}'$, such that s' is a ranking strategy profile and the continuation strategy of s conditional on $\hat{h}$ is identical to the continuation strategy of s' conditional on $\hat{h}'$.

Let S^* be the subset of S^{I} that satisfies the agreement-external condition.[12] The main result of this section shows that payoffs attained by equilibria in S^* are the same as those attained by contractual equilibria.

Theorem 6. *Consider any strategy profile $s \in S^*$ and let V be the set of continuation values supported by s for all histories in H. Then $V \subset V^*$. Furthermore, for any $v \in V^*$, there is a strategy profile in S^* that supports this value from the beginning of the game.*

6 Generalization

In this section, we extend our analysis to a general model with the following features: (1) $n \geq 2$ players, (2) imperfect public monitoring, and (3) heterogeneous discount factors. This extension complements the results of Fong and Surti (2009), who study the Pareto frontier of attainable payoffs in the perfect-monitoring Prisoners' Dilemma with transfers and heterogeneous discount factors, but restrict attention to subgame perfection and strong renegotiation proofness (Farrell and Maskin 1989).[13]

We assume that in the negotiation stage each player obtains a share of the surplus, where the shares are fixed exogenously. A game in this class is defined by:

[12] We could have defined no-fault disagreement and the agreement-internal condition over the set S and then taken the intersection of these and S^{C} to form S^{I}. This produces the equivalent set. It is important, however, that the agreement-external condition be defined relative to S^{I}. That is, the other conditions take precedence over, and are necessary for establishing the meaning of, the agreement-external condition.

[13] Lehrer and Pauzner (1999) also study repeated games with heterogeneous discount factors, but without considering renegotiation.

- A finite number of players n;
- A stage game, featuring a set of action profiles $A = A_1 \times A_2 \times \cdots \times A_n$, a set of public signals S, a signal distribution function $p : A \to \Delta S$, and payoff functions $u_1 : A_1 \times S \to \mathbb{R}, u_2 : A_2 \times S \to \mathbb{R}, \ldots, \hat{u}_n : A_n \times S \to \mathbb{R}$;
- A vector of discount factors $\delta \in [0,1)^n$, where δ_i denotes player i's discount factor; and
- Bargaining weights $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$, with $\pi \geq 0$ and $\sum_{i=1}^n \pi_i = 1$.

In the stage game, the players simultaneously select actions $a_1 \in A_1, a_2 \in A_2, \ldots, a_n \in A_n$, yielding action profile a. Public signal s is then realized according to the probability measure $p(\cdot|a)$. The players mutually observe s but they do not observe each other's actions. We write $u(a,s) = \big(u_i(a_i,s)\big)_{i=1}^n$, and extend p and u to the space of mixed actions.

Since the players' shared information is the history of public signals, and because nothing that occurred in previous periods is payoff-relevant for the future, we assume that the disagreement point in the negotiation phase is commonly known, and therefore conditioned on only the public signals—not on individual actions in the stage game. Further, we suppose that, in their individual actions, the players condition on the history through only the public signals realized in previous periods. Thus, we examine a "perfect public" version of contractual equilibrium, along the lines of perfect public equilibrium.

6.1 Recursive construction

Since utility in average terms is no longer necessarily transferable, we express continuation payoffs in total terms rather than average terms. To make the difference clear in notation, values that were in average terms in previous sections but are now in total terms are shown with a tilde. The construction of a continuation value $\tilde{y}$ from the negotiation phase (Eq. 2 in the basic model) must now represent (i) continuation values as a function of the public signal and (ii) the players' possibly different discount factors. For any two vectors $x, x' \in \mathbb{R}^n$, define

$$x{*}x' \equiv (x_1 x'_1, x_2 x'_2, \ldots, x_n x'_n), \tag{25}$$

which is the vector resulting from component-by-component multiplication. We use the standard notation $x \cdot x'$ for the dot product of x and x'. Also, let $\vec{1}$ be the vector of ones. Then an agreement continuation value y is constructed as follows:

$$\tilde{y} = m + \int_{s \in S} \left[u(\alpha, s) + \delta * \tilde{g}(s) \right] dp(s|\alpha), \tag{26}$$

where $m \in \mathbb{R}_0^n$ is the immediate transfer, $\alpha \in \Delta^{\mathrm{U}} A$ is a mixed action profile to be played in the current period, $\tilde{g}$ gives the continuation value from the start of the next period as a function of the public signal in the current period, and $p(s|\alpha)$ is the probability measure on S given α.[14] A disagreement value is given by:

$$\underline{\tilde{y}} = \int_{s \in S} \left[u(\alpha, s) + \delta * \tilde{g}(s) \right] dp(s|\alpha). \tag{27}$$

Extending Definition 1, we say that $\tilde{g} : S \to \tilde{V}$ enforces α if α is a Nash equilibrium of the game with action-profile space A where, for each $a \in A$, the payoffs are given by

$$\int_{s \in S} \left[u(a, s) + \delta * \tilde{g}(s) \right] dp(s|a). \tag{28}$$

Operators C and D are revised as follows:

$$C(\tilde{V}) \equiv \left\{ \tilde{y} \in \mathbb{R}^2 \,\middle|\, \begin{array}{l} \exists\ m \in \mathbb{R}_0^2,\ \tilde{g} : S \to \tilde{V}, \text{ and } \alpha \in \Delta^{\mathrm{U}} A \\ \text{s.t. } \tilde{g} \text{ enforces } \alpha \text{ and Eq. 26 holds} \end{array} \right\}, \tag{29}$$

$$D(\tilde{V}) \equiv \left\{ \underline{\tilde{y}} \in \mathbb{R}^2 \,\middle|\, \begin{array}{l} \exists\ \tilde{g} : S \to \tilde{V} \text{ and } \alpha \in \Delta^{\mathrm{U}} A \\ \text{s.t. } \tilde{g} \text{ enforces } \alpha \text{ and Eq. 27 holds} \end{array} \right\}. \tag{30}$$

The function B requires no modification for the general model. As before, $\tilde{V}$ is said to satisfy negotiated self-generation (NSG) if $\tilde{V} = \operatorname{co} B(C(\tilde{V}), D(\tilde{V}))$. The definitions of dominance and contractual equilibrium are the same as before.

With A and S finite, we can guarantee existence.

Theorem 7. *Consider any n-player repeated game with imperfect public monitoring, heterogenous discount factors, renegotiation, and transfers — defined by $\langle A, S, p, u, \delta, \pi \rangle$. If A*

[14] Recall that $\mathbb{R}_0^n \equiv \left\{ m' \in \mathbb{R}^n \mid \sum_{i=1}^n m'_i = 0 \right\}$ is the set of balanced transfers.

and S are finite and $\delta_i \in [0,1)$ for all i, then the game has a unique CEV set $\tilde{V}^*$.

The proof examines a fixed-point problem for a transformation of the self-generation operator $\operatorname{co} B(C(\cdot), D(\cdot))$. The transformation normalizes the set of continuation values by subtracting the welfare level, in a direction that accounts for the heterogeneous discount factors and bargaining shares. Then we can apply a fixed point argument to the normalized set of continuation values—which is a hyperpolygon in $\mathbb{R}^n$ rather than a line segment in $\mathbb{R}^2$. Since the normalized bargaining operator preserves compactness, and is bounded, monotone, and continuous on decreasing sets, it has a largest fixed point.

6.2 Two-player games

In this subsection, we give a characterization of the CEV set along the lines of the construction in Section 4.1 for the special case in which there are two players (with imperfect public monitoring and where the players have possibly different discount factors $\delta = (\delta_1, \delta_2)$). We assume S is finite. We first work through the construction of z^2, which is player 1's most preferred point in the CEV set. In this environment, Eq. 11 becomes:

$$\tilde{z}_1^2 = \max_{\underline{\tilde{y}}, \tilde{g}, \alpha} \; \pi_2 \underline{\tilde{y}}_1 - \pi_1 \underline{\tilde{y}}_2 + \pi_1 \tilde{L}, \qquad \text{s.t.} \begin{cases} \underline{\tilde{y}} = \sum_s p(s|\alpha)\big[u(\alpha, s) + \delta * \tilde{g}(s)\big], \\ \tilde{g} : S \to \tilde{V} \text{ enforces } \alpha. \end{cases} \tag{31}$$

Define $\eta(s) \equiv \tilde{g}_1(s) - \tilde{z}_1^2$ and $\hat{u}_i(\alpha) \equiv \sum_s p(s|\alpha) u_i(\alpha_i, s)$. We shall also write $\hat{\eta}(\alpha) \equiv \sum_s p(s|\alpha)\eta(s)$. After rearranging terms as before, the optimization problem of Eq. 12 becomes:

$$\tilde{z}_1^2 = \max_{\eta, \alpha} \; \big(\pi_2 \hat{u}_1(\alpha) - \pi_1 \hat{u}_2(\alpha)\big) + \big(\pi_2 \delta_1 \tilde{z}_1^2 - \pi_1 \delta_2 \tilde{z}_2^2\big) + \big(\pi_2 \delta_1 + \pi_1 \delta_2\big)\hat{\eta}(\alpha) + \pi_1 \tilde{L}, \qquad \text{s.t.} \begin{cases} \eta : S \to [\tilde{z}_1^1 - \tilde{z}_1^2, 0], \text{ with } \hat{\eta}(\alpha) \equiv \sum_s p(s|\alpha)\eta(s), \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, \hat{u} + \delta * \big((\hat{\eta}, -\hat{\eta}) + \tilde{z}^2\big)\big\rangle. \end{cases} \tag{32}$$

Letting $\psi \equiv \pi_1\delta_2 + \pi_2\delta_1$, we can now redefine the function $\overline{\gamma}$ by

$$\begin{aligned} \overline{\gamma}(\tilde{d}) \equiv \max_{\eta,\alpha} \;& \frac{\pi_2}{1-\psi}\hat{u}_1(\alpha) - \frac{\pi_1}{1-\psi}\hat{u}_2(\alpha) + \frac{\psi}{1-\psi}\hat{\eta}(\alpha), \\ \text{s.t.} \;& \begin{cases} \eta : S \to [-\tilde{d}, 0], \text{ with } \hat{\eta}(\alpha) \equiv \sum_s p(s|\alpha)\eta(s), \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, \hat{u} + \delta * (\hat{\eta}, -\hat{\eta})\big\rangle, \end{cases} \end{aligned} \tag{33}$$

We thus have the following partial characterization of $\tilde{z}^2$, the general version of Eq. 13:

$$\tilde{z}_1^2 = \frac{1-\delta_2}{1-\psi}\pi_1\big(\tilde{z}_1^2 + \tilde{z}_2^2\big) + \overline{\gamma}(\tilde{z}_1^2 - \tilde{z}_1^1). \tag{34}$$

By similar calculations, we obtain a partial characterization of $\tilde{z}^1$:

$$\tilde{z}_1^1 = \frac{1-\delta_2}{1-\psi}\pi_1\big(\tilde{z}_1^1 + \tilde{z}_2^1\big) + \underline{\gamma}(\tilde{z}_1^2 - \tilde{z}_1^1). \tag{35}$$

where $\underline{\gamma}$ is defined by

$$\begin{aligned} \underline{\gamma}(\tilde{d}) \equiv \min_{\eta,\alpha} \;& \frac{\pi_2}{1-\psi}\hat{u}_1(\alpha) - \frac{\pi_1}{1-\psi}\hat{u}_2(\alpha) + \frac{\psi}{1-\psi}\hat{\eta}(\alpha), \\ \text{s.t.} \;& \begin{cases} \eta : A \to [0, \tilde{d}], \text{ with } \hat{\eta}(\alpha) \equiv \sum_s p(s|\alpha)\eta(s), \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, \hat{u} + \delta * (\hat{\eta}, -\hat{\eta})\big\rangle. \end{cases} \end{aligned} \tag{36}$$

To calculate the CEV set, we find the maximal fixed point of $\Gamma = \overline{\gamma} - \underline{\gamma}$, which yields the payoff span $\tilde{d}^*$. We then have to calculate the level of the CEV set, which is a bit more involved than in the basic model because it is the infinite sum of discounted payoffs where the players have different discount factors. Note that the level satisfies:

$$\begin{aligned} \tilde{L}^* = \tilde{z}_1^{1*} + \tilde{z}_2^{1*} = \max_{\tilde{g},\alpha} \;& \hat{u}_1(\alpha) + \hat{u}_2(\alpha) + \sum_s p(s|\alpha)\delta \cdot \tilde{g}(s), \\ \text{s.t. } & \tilde{g} : S \to \tilde{V}^* \text{ enforces } \alpha, \end{aligned} \tag{37}$$

because the objective function here is the joint value from the current period and the constraint requires that continuation values in the following period be chosen from $\tilde{V}^*$.

With the same steps taken above, we rewrite this maximization problem by substituting for $\tilde{g}$ using the function η, where we have $\tilde{g}_1(s) = \eta(s) + \tilde{z}_1^1 + \tilde{d}^*$ and $\tilde{g}_2(s) = \tilde{z}_2^1 - \tilde{d}^* - \eta(s)$.

This yields:

$$\tilde{z}_1^{1*} + \tilde{z}_2^{1*} = \delta_1 \tilde{z}_1^{1*} + \delta_2 \tilde{z}_2^{1*} + \rho(\tilde{d}^*) \tag{38}$$

where

$$\begin{aligned}\rho(\tilde{d}) \equiv \max_{\eta,\alpha} \;\; & \hat{u}_1(\alpha) + \hat{u}_2(\alpha) + (\delta_1 - \delta_2)\hat{\eta}(\alpha) + (\delta_1 - \delta_2)\tilde{d}, \\ \text{s.t.} \;\; & \begin{cases} \eta : S \to [0, \tilde{d}], \text{ with } \hat{\eta}(\alpha) \equiv \sum_s p(s|\alpha)\eta(s), \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, \hat{u} + \delta * (\hat{\eta}, -\hat{\eta}) \big\rangle. \end{cases}\end{aligned} \tag{39}$$

Noting that $\tilde{d}^* = \tilde{z}_1^{2*} - \tilde{z}_1^{1*}$, we can rewrite Eq. 35 for the CEV payoff span $\tilde{d}^*$ as

$$\tilde{z}_1^{1*} = \frac{1 - \delta_2}{1 - \psi} \pi_1 \big(\tilde{z}_1^{1*} + \tilde{z}_2^{1*} \big) + \underline{\gamma}(\tilde{d}^*). \tag{40}$$

Eq. 38 and Eq. 40 form a system of equations with unknowns $\tilde{z}_1^{1*}$ and $\tilde{z}_2^{1*}$. The solution is:

$$\tilde{z}_1^{1*} = \frac{1}{1 - \delta_1} \big[\pi_1 \rho(\tilde{d}^*) + (1 - \psi)\underline{\gamma}(\tilde{d}^*) \big] \tag{41}$$

$$\tilde{z}_2^{1*} = \frac{1}{1 - \delta_2} \big[\pi_2 \rho(\tilde{d}^*) - (1 - \psi)\underline{\gamma}(\tilde{d}^*) \big]. \tag{42}$$

We then have $\tilde{z}^{2*} = \tilde{z}^{1*} + (\tilde{d}^*, -\tilde{d}^*)$ and $\tilde{L}^* = \tilde{z}_1^{1*} + \tilde{z}_2^{1*} = \tilde{z}_1^{2*} + \tilde{z}_2^{2*}$. In the special case of equal discount factors $\delta_1 = \delta_2 \equiv \delta_0$, ρ simplifies and we have $\tilde{L}^* = \rho(\tilde{d}^*)/(1 - \delta_0)$.

7 Conclusion

We have introduced *contractual equilibrium*, a new approach to modeling negotiation in repeated games. Contractual equilibrium allows disagreement play to arise endogenously, in a way consistent with a well-defined extensive form. We have provided a complete characterization of contractual equilibria in repeated games with monetary transfers, allowing for imperfect monitoring and heterogeneous discount factors. Contractual equilibrium exists, yields a unique welfare level, and is tractable in applications.

The underlying principles of our approach are not restricted to repeated games. In dynamic games, where a payoff-relevant state can vary over time and shift endogenously in response to players' actions, the same notions of agreement and disagreement apply. Endogenous disagreement can be particularly powerful in this setting, since the actions taken under disagreement can change the set of feasible continuation values.

The legal environment—the set of enforceable contracts and the enforcement technology—can introduce dynamics if long-term contracts are enforceable. Then the players' cooperative action of signing an agreement, though it has no effect on the feasibility of the equilibrium path due to renegotiation, can leverage the threat of enforcement to change the feasible set of continuation values after future disagreements and deviations.

Transferable utility is important to our results on existence and uniqueness, as well as our characterizations. Contractual equilibrium can still be applied to environments without transfers, but if the set of feasible continuation values is identical under agreement and disagreement then every contractual equilibrium can be made strongly renegotiation proof in the sense of Farrell and Maskin (1989). Contractual equilibrium differs substantively from strong renegotiation proofness whenever the set of feasible payoffs is larger under agreement than under disagreement—such as when transfers are possible but costly, or where agreement is instantiated in a legal contract.

We showed how contractual equilibrium applies to games with more than two players, using exogenously given bargaining shares—such as could be derived from an n-player extension of the Nash bargaining solution. To use the Nash bargaining solution, we need a well-defined disagreement point, which is why we impose the refinement of no-fault disagreement on equilibria in our fully noncooperative model. To use other bargaining notions, such as the Myerson (1977) value or core-based notions that accommodate externalities, we could relax no-fault disagreement to allow for arbitrary coalition structures to arise under disagreement, with endogenous continuation play for each coalition structure.

Appendix

Proof of Theorem 2: To show the dependence of $\overline{\gamma}$, $\underline{\gamma}$, and Γ on δ, let us write $\overline{\gamma}^\delta$, $\underline{\gamma}^\delta$, and $\Gamma^\delta(d)$.

Claim 1. *Given any $\delta, \delta' \in (0,1)$, and $d \geq 0$, let $d' \equiv \frac{\delta(1-\delta')}{\delta'(1-\delta)}d$. Then $\Gamma^{\delta'}(d') = \Gamma^\delta(d)$.*

Proof. Consider the definition of $\overline{\gamma}$ (see Eq. 14) and suppose that α and η satisfy the constraints for discount factor δ. Let $\eta' \equiv \frac{\delta(1-\delta')}{\delta'(1-\delta)}\eta$. Note that the first constraint is equivalent to $\eta' \in [-d', 0]$. Regarding the second constraint, observe that

$$(1-\delta)u + \delta(\eta, -\eta) = \frac{1-\delta}{1-\delta'}\left[(1-\delta')u + \delta'(\eta', -\eta')\right]. \tag{43}$$

Thus, α and η satisfy the second constraint for δ if and only α and η' satisfy the same constraint for δ'. Finally, note that the value of the objective function of Eq. 14 at α, η, δ is equal to the value at α, η', δ'. These facts imply that $\overline{\gamma}^{\delta}(d) = \overline{\gamma}^{\delta'}(d')$. By the same reasoning, we have $\underline{\gamma}^{\delta}(d) = \underline{\gamma}^{\delta'}(d')$, and so $\Gamma^{\delta}(d) = \Gamma^{\delta'}(d')$. □

Fix δ. If $\Gamma^{\delta}(\infty) > 0$ (which is equivalent to $\overline{\gamma}^{\delta}(\infty) > \underline{\gamma}^{\delta}(\infty)$), then there exists a number $\hat{d} > 0$ such that $\Gamma^{\delta}(\hat{d}) > 0$. Let $\varepsilon \equiv \Gamma^{\delta}(\hat{d})$ and, for any δ', let $d' \equiv \frac{\delta(1-\delta')}{\delta'(1-\delta)}\hat{d}$. From Claim 1 we see that $\Gamma^{\delta'}(d') = \varepsilon$. As δ' converges to one, d' converges to zero, which implies that the maximal fixed point of Γ is bounded below by ε for sufficiently large discount factors.

Finally, it is clear from the definition of L^* (see Eq. 18) that, if d^* is bounded away from zero for large discount factors, then any stage-game action profile can be supported for δ sufficiently large. This proves the first statement of the theorem. Regarding the second statement, observe that $\overline{\gamma}^{\delta}(\infty) > \underline{\gamma}^{\delta}(\infty)$ implies that $\Gamma^{\delta}(d) = 0$ for all δ and d. Thus, the maximal fixed point of Γ is zero regardless of δ and only stage-game Nash equilibrium action profiles can be supported. ■

Proof of Theorem 3: Let $\hat{a} \equiv \arg\max_a \sum_i u_i(a)$. Under these strategies, let $\hat{z}^i = u(a^i) + \pi\big(\sum_j \big(u_j(\hat{a}) - u_j(a^i)\big)$. Note that $\hat{z}^i$ does not depend on δ. Since $(\pi_2, -\pi_1) \cdot \big(u(a^2) - u(a^1)\big) > 0$, the payoff span $\hat{d} = \hat{z}_1^2 - \hat{z}_1^1$ is strictly positive and constant in δ.

We must check that the sequential rationality constraints are satisfied. Under disagreement in state ω^i, player i is playing a stage game best response, and anticipates remaining in state ω^i regardless of her action. Under agreement in either state, player i anticipates a loss of $\frac{\delta}{1-\delta}\frac{1}{2}\hat{d}$ if she deviates. Similarly, under disagreement in state ω^{-i}, player i anticipates a loss of $\frac{\delta}{1-\delta}\hat{d}$ if she deviates. Hence all actions are sequentially rational for δ sufficiently high. Note that the payoff span of the equilibrium we construct is not necessarily the full payoff span of V^*. ■

Proof of Theorem 4: It suffices to restrict attention to the stage game and find α^1 and α^2 as described in Theorem 3. Choose α^*, η, and k satisfying the suppositions. For any $\varepsilon > 0$ let $\alpha_i^{-i}(\varepsilon) \equiv \alpha_i^* + \varepsilon$ and $\alpha_i^i(\varepsilon) \equiv BR_i\big(\alpha_{-i}^i(\varepsilon)\big)$. Then for all $\varepsilon < \eta/2k$, $\alpha_i^* - \varepsilon k < \alpha_i^i(\varepsilon) < \alpha_i^* + \varepsilon k$.

Near α^*, since the utility functions are differentiable and α^* is an equilibrium, for ε small and α_{-i} sufficiently close to α_{-i}^* it follows that $\big|u_i\big(\alpha_i^i(\varepsilon), \alpha_{-i}\big) - u_i\big(\alpha_i^{-i}(\varepsilon), \alpha_{-i}\big)\big|$

is on the order of at most ε^2. Since $du_i / d\alpha_{-i} < -1/k$, for α_i sufficiently close to α_i^* it also follows that $u_i\big(\alpha_i, \alpha_{-i}^{-i}(\varepsilon)\big) - u_i\big(\alpha_i, \alpha_{-i}^{i}(\varepsilon)\big) > 0$ is on the order of at least ε. Hence, for $\varepsilon > 0$ sufficiently small, each player i strictly prefers α^{-i} to α^i. Since player i is best responding at α^i, the conditions of Theorem 3 are satisfied. ■

Proof of Theorem 5: Define

$$\hat{\Gamma}(d, \alpha^1, \alpha^2) \equiv \max_{\eta^1, \eta^2} \; \pi_2\left(u_1(\alpha^1) - u_1(\alpha^2)\right) - \pi_1\left(u_2(\alpha^1) - u_2(\alpha^2)\right) + \frac{\delta}{1-\delta}\big(\eta^1(\alpha^1) - \eta^2(\alpha^2)\big),$$
$$\text{s.t.} \begin{cases} \eta^1 : A \to \big[-d, 0\big], \text{ extended to } \Delta^{\mathrm{U}} A, \\ \eta^2 : A \to \big[0, d\big], \text{ extended to } \Delta^{\mathrm{U}} A, \\ \alpha^1 \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, (1-\delta)u + \delta(\eta^1, -\eta^1)\big\rangle, \\ \alpha^2 \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, (1-\delta)u + \delta(\eta^2, -\eta^2)\big\rangle. \end{cases} \tag{44}$$

That is, $\Gamma(d) = \max_{\alpha^1, \alpha^2} \hat{\Gamma}(d, \alpha^1, \alpha^2)$. Observe that the argmax (η^1, η^2) of $\hat{\Gamma}(d, \alpha^1, \alpha^2)$ is independent of π_1 and π_2. Hence $\hat{\Gamma}(d, \alpha^1, \alpha^2)$ is maximized at either $\pi = (0, 1)$ or $\pi = (1, 0)$, as is $\Gamma(d)$, and therefore $\max_d \{d : d = \Gamma(d)\}$. ■

Proof of Theorem 6: First we construct a subgame perfect equilibrium in the fully noncooperative game whose equilibrium path continuation values are all contained in V^*, and show that the equilibrium is a member of S^{I}. Then we show that every equilibrium in S^{I} is dominated by the one that we constructed, and therefore every equilibrium in S^* attains the same welfare level as V^*.

Step 1: An equilibrium that attains V^*. Recall that $V^* = \mathrm{co}\{z^1, z^2\}$, where z^1 is the endpoint that favors player 2 and z^2 is the endpoint that favors player 1, and that $L^* = z_1^1 + z_2^1$ is the welfare level of V^*.

Consider z^1. Because V^* is NSG, we know that $L^* = max_{y \in C(V^*)} y_1 + y_2$ and there exists a point $\underline{y}^1 \in D(V^*)$ such that $z^1 = \underline{y}^1 + \pi(L^* - \underline{y}_1^1 - \underline{y}_2^1)$. From the definition of D (recall Eq. 5), we know there exists a function $g^1 : A \to V^*$ and an action profile $\alpha^1 \in \Delta^{\mathrm{U}} A$ such that g^1 enforces α^1 and $\underline{y}^1 = (1-\delta)u(\alpha^1) + \delta g^1(\alpha^1)$. Note that g can be implemented by randomizing over z^1 and z^2 using the public randomization device. A similar construction works for z^2.

We construct a strategy profile s^* that, at the negotiation phase of each period, differentiates between two states, 1 and 2. The implied disagreement point will be $\underline{y}^1$ in state 1, whereas it will be $\underline{y}^2$ in state 2. In state $k = 1, 2$, if player i is the offerer and player j is the responder, then s^* prescribes that player i offer $\lambda = (w^*, m^{ki})$ where m^{ki} is defined so that $w_j^* + m_j^{ki} = \underline{y}_j^k$. As for player j's response more generally, if player i offers $\lambda = (w, m)$ where $w_j + m_j \geq \underline{y}_j^k$ and $w \in D(V^*)$ then player j should respond "yes," and they should play $\tilde{\alpha}^w$ and use the public randomization device to achieve $\tilde{g}^w$ (as a function of a). If player i makes any other offer or player j says "no" then the players are to play $\alpha^k \in \Delta^{\mathrm{U}} A$ with mixing over the states in the next period to achieve g^k. An arbitrary mixture between the two states can occur in the first period.

By construction, s^* is a subgame-perfect equilibrium that satisfies meaningful communication and no-fault disagreement. In particular, note that $\lambda = (w^*, m^{ki})$ is player i's optimal offer in state k given player j's acceptance rule and the way the players coordinate in the continuation game as a function of the offer and response. Observe that in state k, the expected payoff vector is z^k. For every $w \in D(V^*)$ there is a continuation from the action phase in which the players obtain w; this occurs when the offerer states $\lambda = (w, m)$ and the responder says "yes." Further, for every $w \notin D(V^*)$ there is *no* continuation from the action phase in which the players obtain w.

Moreover, after every history the offerer has the option to propose any continuation payoff in $D(V^*)$, and knows that such a proposal will be accepted if it provides a transfer that makes the responder indifferent between her disagreement payoff and her proposed payoff. This implies that s^* is an element of S^{I}. By construction, its set of continuation values following histories in H is contained in V^*, so there is an appropriate randomization to attain any value in V^* at the beginning of the game.

Step 2: Characterizing Equilibria in $s \in S^{\mathrm{I}}$. The next step in the proof is to establish some properties shared by every strategy profile in S^{I}. Consider any $s \in S^{\mathrm{I}}$. First observe that there must exist a level $L \in \mathbb{R}$ such that, for every history in H, the continuation value v satisfies $v_1 + v_2 = L$. This follows from the no-fault disagreement and agreement-internal conditions.

Suppose, for instance, that two levels are supported; that is, there is a history $h \in H$ from which v is the continuation value, and there is another history $h' \in H$ from which v' is the continuation value, with $v_1 + v_2 < v_1' + v_2'$. We shall derive a contradiction. There is a history $\hat{h}' = (h'; \omega, k, \lambda, m) \in \hat{H}$ from which the continuation payoff vector w'' satisfies

$v_1' + v_2' \leq w_1'' + w_2''$. Further, following h there must be a realization ω from which the continuation value y satisfies $y_1 + y_2 \leq v_1 + v_2$. But then there exists a transfer m'' such that $m'' + w''$ stricly exceeds y (for both players) and also exceeds the disagreement point in force from $(h; \omega)$. Suppose that following $(h; \omega)$ the offerer states $\lambda = (w'', m'')$. By the agreement-internal and no-fault disagreement conditions, the responder rationally must say "yes," leading to the continuation value $m'' + w''$. Because the offerer strictly prefers the continuation value $m'' + w''$ to v, he strictly prefers to deviate.

The second observation to make is that for any strategy in S^{I}, the set of continuation values (from histories in H) satisfies a negotiated self-generation condition that is related to the one developed in Section 3.2. For any set $\underline{Y} \subset \mathbb{R}^2$ and any level $L \in \mathbb{R}$, define:

$$\tilde{B}(L, \underline{Y}) \equiv \{\underline{y} + \pi(L - \underline{y}_1 - \underline{y}_2) \mid \underline{y} \in \underline{Y} \text{ satisfying } \underline{y}_1 + \underline{y}_2 \leq L\}. \tag{45}$$

Consider any $s \in S^{\mathrm{I}}$ and let V be its set of continuation values over all histories in H. It is the case that, for some L that is supported by using continuation values in $D(V)$, we have $V \subset \tilde{B}(L, D(V))$. The reason for this is that the disagreement condition implies that, for each history $(h; \omega)$ (with $h \in H$), there is a disagreement point $\underline{y} \in D(V)$. By the agreement-internal condition, in the equilibrium continuation when player i is the offerer, the continuation payoff is $\underline{y}_i + (L - \underline{y}_1 - \underline{y}_2)$ for player i and $\underline{y}_j$ for the other player j. This is because the offerer can suggest a continuation payoff yielding joint value L that the responder must accept and that gives the offerer arbitrarily close to the full surplus over the disagreement point.

Comparing $\tilde{B}$ and B, it is clear that the span of V cannot exceed the span of V^*, so $L \leq L^*$. Furthermore, V^* dominates V. This means that s^* constructed above is a ranking strategy profile and no other ranking strategy profile supports continuation values from the action phase that are outside the set $D(V^*)$. This implies that $s^* \in S^*$. By the agreement-external condition, every strategy profile in S^* achieves the level L^* from every history in H. Otherwise, there would be a strictly profitable deviation for an offerer to suggest some $\lambda = (w, m)$ satisfying $w_1 + w_2 = L^*$ and which the responder strictly wants to accept. These facts are sufficient to prove the theorem. ■

Proof of Theorem 7: Let $\mathcal{X}$ denote the set of compact subsets of $\mathbb{R}^n$, and let $\mathcal{X}_0$ denote the set of compact subsets of $\mathbb{R}^n_0$. For any set $X \in \mathcal{X}$ and any point $x' \in \mathbb{R}^n$, let the sum be defined by $X + x' \equiv \{x + x' \mid x \in X\}$. For every $X \in \mathcal{X}$, define $b(X) \equiv \operatorname{co} B(C(X), D(X))$,

$L(X) \equiv \max_{x \in X} \sum_{i=1}^{n} x_i$, and $\beta(X) \equiv b(X) - \pi L(b(X))$. The definition of L is consistent with its use in Section 4, but we now write it in relation to an arbitrary set. The function β normalizes the output of b so that points in the resulting sets have a joint value of zero. The angle by which this normalization takes place (in the direction of π) is critical for the analysis below.

Claim 2. *Function b maps $\mathcal{X}$ to $\mathcal{X}$, and function β maps $\mathcal{X}$ to $\mathcal{X}_0$. For every $\tilde{v} \in b(X)$, it is the case that $\sum_{j=1}^{n} \tilde{v}_j = L(b(X))$. Furthermore, for any $X \in \mathcal{X}$ and $x \in \mathbb{R}^n$, $b(X + x) = b(X) + \delta * x$ and $L(X + x) = L(X) + L(\{x\})$.*

Proof. By upper hemi-continuity of the Nash equilibrium correspondence, the operators C and D preserve compactness. The Nash bargaining solution is continuous in the disagreement point and maximal joint value when applied to bargaining problems with transferrable utility. Furthermore, $\beta(X)$ is a linear transformation of $b(X)$ that makes every point balanced (joint value of zero). These facts imply the first part of the claim. The second part follows from transferrable utility and the bargaining solution (as in Lemma 1). Regarding the third part, note that adding x to every point in X merely shifts the set of continuation values from the next period. Referring to Eq. 26, this is equivalent to replacing $\tilde{g}(s)$ with $\tilde{g}(s) + x$. Thus, the set of values that can be supported from the current period ($\tilde{y}$ vectors) uniformly shifts by $\delta * x$. Consequently, the set of bargaining outcomes likewise shifts. Finally, the level clearly changes as indicated. $\square$

Hereinafter, all sets that we consider are understood to be compact subsets of $\mathbb{R}$. To establish the existence of a unique CEV set, we first characterize the NSG sets—the fixed points of b. To identify and compare NSG sets, we shall work with the function β. We start by demonstrating a relation between the fixed points of b and β. For each player i define $\phi_i \equiv \pi_i/(1 - \delta_i)$, write $\phi = (\phi_1, \phi_2, \ldots, \phi_n)$, and let $\Phi \equiv \sum_{j=1}^{n} \phi_j$. Claim 3 provides a linear relationship between the fixed points of b and β, in the direction ϕ.[15]

Claim 3. *If $\tilde{V} \in \mathcal{X}$ and $X = \tilde{V} - \frac{\phi}{\Phi} L(\tilde{V})$, then $\tilde{V} = b(\tilde{V})$ implies $X = \beta(X)$. If $X \in \mathcal{X}_0$ and $\tilde{V} = X + \phi L(b(X))$, then $X = \beta(X)$ implies $\tilde{V} = b(\tilde{V})$.*

[15] This relationship does not hold for other sets in general.

Proof. To prove the first part of the claim, we start with the following algebraic steps:

$$
\begin{aligned}
\beta(X) &= b(X) - \pi L(b(X)) \\
&= \tilde{V} - \delta * \frac{\phi}{\Phi} L(\tilde{V}) - \pi L(\tilde{V} - \delta * \frac{\phi}{\Phi} L(\tilde{V})) \\
&= \tilde{V} - \delta * \frac{\phi}{\Phi} L(\tilde{V}) - \pi L(\tilde{V}) \Big(1 - \delta \cdot \frac{\phi}{\Phi}\Big) = \tilde{V} - \Big(\pi + (\vec{1} - \pi)\delta * \frac{\phi}{\Phi}\Big) L(\tilde{V}).
\end{aligned} \tag{46}
$$

The second line uses the property of b from Claim 2 and that $V = b(V)$. The third line uses the property of L from Claim 2. Note that

$$
\Phi - \delta \cdot \phi = \sum_{j=1}^{n} \phi_j (1 - \delta_j) = \sum_{j=1}^{n} \frac{\pi_j}{1 - \delta_j} (1 - \delta_j) = 1. \tag{47}
$$

Thus,

$$
\beta(X) = \tilde{V} - \delta * \frac{\phi}{\Phi} L(\tilde{V}) - \pi \frac{1}{\Phi} L(\tilde{V}) = \tilde{V} - (\delta * \phi + \pi) \frac{1}{\Phi} L(\tilde{V}). \tag{48}
$$

It can be verified that $\delta_i \phi_i + \pi_i = \phi_i$, which means that $\delta * \phi + \pi = \phi$. Therefore $\beta(X) = \tilde{V} - \frac{\phi}{\Phi} L(\tilde{V}) = X$.

To prove the second part of the claim, we perform the following algebraic steps:

$$
\begin{aligned}
b(\tilde{V}) &= b(X + \phi L(b(X)) \\
&= b(X) + \delta * \phi L(b(X)) = b(X) - \pi L(b(X)) + \pi L(b(X)) + \delta * \phi L(b(X)) \\
&= \beta(X) + (\pi + \delta * \phi) L(b(X)) = X + \phi L(b(X)) = \tilde{V}.
\end{aligned} \tag{49}
$$

The second line uses the property of b from Claim 2 and the third line uses the definition of β. The fourth line uses the assumption that $\beta(X) = X$ and that $\delta * \phi + \pi = \phi$, which we showed above. □

We next show by construction that β has a dominant fixed point. We start by identifying some properties of β.

Claim 4. *The function β is* **monotone***: for every $X, X' \in \mathcal{X}$, $X \subset X'$ implies $\beta(X) \subset \beta(X')$. Furthermore, β is* **continuous on decreasing sequences***: for every sequence $\{X^k\} \subset \mathcal{X}$ with $X^{k+1} \subset X^k$ for all k, if X^k converges to X in the Hausdorff metric then $\beta(X^k)$ converges to $\beta(X)$.*

Proof. To prove the first part of the claim, take sets $X, X' \in \mathcal{X}$ such that $X \subset X'$. Clearly D is monotone so $D(X) \subset D(X')$. By construction, we know that for each point $\tilde{y} \in B(C(X), D(X))$ there is an element $\underline{\tilde{y}} \in D(X)$ such that $\tilde{y} = \underline{\tilde{y}} + \pi[L(D(X)) - \vec{1}\cdot\underline{\tilde{y}}]$. Using the same disagreement point and choosing $\tilde{y}' = \underline{\tilde{y}} + \pi[L(D(X')) - \vec{1}\cdot\underline{\tilde{y}}]$, we have that $\tilde{y}' \in B(C(X'), D(X'))$. Thus, $\tilde{y} \in b(X)$, $\tilde{y}' \in b(X')$, and these points lie on the same line in direction π. Recall that β normalizes by subtracting a vector proportional to π, so that resulting points have joint value of zero. This means that $\tilde{y} - \pi L(b(X) = \tilde{y}' - \pi L(b(X'))$, which proves $\beta(X) \subset \beta(X')$.

To prove the second part of the claim, consider any decreasing sequence $\{X^k\} \subset \mathcal{X}$ such that X^k converges to X for some $X \in \mathcal{X}$. Because the Nash equilibrium correspondence is upper hemi-continuous, we know that $\lim_{k\to\infty} D(X^k) \subset D(X)$. Since $\{X^k\}$ is decreasing, we have that $X \subset X^k$ for all k. Because D is monotone, $D(X) \subset D(X^k)$ holds, and this implies that $D(X) \subset \lim_{k\to\infty} D(X^k)$. Thus, $D(X^k)$ converges to $D(X)$. By the same reasoning, $C(X^k)$ converges to $C(X)$. The function B is continuous as described in the proof of Claim 2. Thus, $b(X^k)$ converges to $b(X)$ and so $\beta(X^k)$ converges to $\beta(X)$. □

The next claim follows from the fact that weakly more action profiles in the stage game can be enforced for larger sets of continuation values.

Claim 5. *For $X, X' \in \mathcal{X}_0$, $X \subset X'$ implies that $L(b(X)) \leq L(b(X'))$.*

To construct a dominant fixed point of β, start with a large element of $\mathcal{X}_0$ that is guaranteed to be a superset of any fixed point. Since S and A are finite, there exists $\zeta \in \mathbb{R}_+$ such that all stage-game payoffs are bounded below by $-\zeta(1 - \max_j \delta_j)$ and above by $\zeta(1 - \max_j \delta_j)$. Let $X^1 \equiv \{x \in \mathbb{R}^n_0 \mid -\zeta \leq x \leq \zeta \text{ for all } i\}$. Then every fixed point of β is a subset of X^1 and that $\beta(X^1) \subset X^1$. Define the sequence $\{X^k\}$ inductively by $X^{k+1} \equiv \beta(X^k)$, for all $k > 1$. Since β is monotone, this sequence is decreasing. Furthermore, $\{X^k\} \subset \mathcal{X}_0$. Since every decreasing sequence of compact sets in a Euclidean space converges, there exists $X^* \in \mathcal{X}_0$ to which X^k converges.

Claim 4 implies that $X^* = \beta(X^*)$. To see this, note that $X^* \subset X^{k+1} = \beta(X^k)$ for all k. Because β is continuous on decreasing sequences, $\beta(X^k)$ converges to $\beta(X^*)$ and so we have that $X^* \subset \beta(X^*)$. In addition, because β is monotone and $X^* \subset X^k$, we have $\beta(X^*) \subset \beta(X^k) = X^{k+1}$ for all k. That X^{k+1} converges to X^* then implies that $\beta(X^*) \subset X^*$.

Next we argue that every fixed point of β is a subset of X^*. Suppose that this were not the case, so that there is a set $\hat{X} \in \mathcal{X}_0$ such that $\hat{X} = \beta(\hat{X})$ but $\hat{X} \not\subset X^*$. Then we can find a positive integer K such that $\hat{X} \subset X^k$ for all $k \leq K$, but $\hat{X} \not\subset X^{K+1}$. This violates monotonicity of β, which requires $\hat{X} = \beta(\hat{X}) \subset \beta(X^K) = X^{K+1}$.

Thus, we have established that X^* is a fixed point of β and every other fixed point of β is contained in X^*. Define $\tilde{V}^* \equiv X^* + \phi L(b(X^*))$. We finish the proof by showing that $\tilde{V}^*$ is an NSG set for the game and it dominates every other NSG set, so it is the CEV set. That $\tilde{V}^*$ is an NSG set follows immediately from Claim 3. For the second step, consider any other NSG set $\tilde{V}$ and define $\tilde{X} \equiv \tilde{V} - \frac{\phi}{\Phi} L(\tilde{V})$. From Claim 3 we know that $\tilde{X}$ is a fixed point of β. We also know that $\tilde{X} \subset X^*$. From the relationship between fixed points of b and β, we know that $\tilde{V} = \tilde{X} + \phi L(b(\tilde{X}))$. Take any $\tilde{v} \in \tilde{V}$ and let $\tilde{x} \in \tilde{X}$ be such that $\tilde{v} = \tilde{x} + rL(b(\tilde{X}))$. Since $\tilde{X} \subset X^*$, we have that $\tilde{x} \in X^*$ and thus $\tilde{v}' \equiv \tilde{x} + \phi L(b(X^*)) \in \tilde{V}$. Comparing $\tilde{v}$ and $\tilde{v}'$, we see that $\tilde{v}' - \tilde{v} = \phi\big(L(b(X^*)) - L(b(\tilde{X}))\big)$. From Claim 5, we know that $L(b(X^*)) \geq L(b(\tilde{X}))$. In addition, note that $\phi_i \geq 0$ for all i. These facts imply that $\tilde{v}' \geq \tilde{v}$ (that is, $\tilde{v}'_i \geq \tilde{v}_i$ for every player i), which proves that $\tilde{V}^*$ dominates $\tilde{V}$. Thus, $\tilde{V}^*$ is the unique CEV set. ■

References

Dilip Abreu and David G. Pearce. A perspective on renegotiation in repeated games. In Reinhard Selten, editor, *Game Equilibrium Models*, volume II, pages 44–55. Springer-Verlag, Berlin, 1991. URL http://www.princeton.edu/~dabreu/index_files/A%20perspective%20on%20renegotiation.pdf.

Dilip Abreu, David G. Pearce, and Ennio Stacchetti. Toward a theory of discounted repeated games with imperfect monitoring. *Econometrica*, 58(5):1041–1063, September 1990.

Dilip Abreu, David G. Pearce, and Ennio Stacchetti. Renegotiation and symmetry in repeated games. *Journal of Economic Theory*, 60:217–240, 1993.

Philippe Aghion and Patrick Bolton. An incomplete contracts approach to financial contracting. *The Review of Economic Studies*, 59(3):473–494, July 1992.

Philippe Aghion and Jean Tirole. Formal and real authority in organizations. *Journal of Political Economy*, 105(1):1–29, February 1997.

Geir G. Asheim. Extending renegotiation-proofness to invinite horizon games. *Games and Economic Behavior*, 3(3):278–294, August 1991.

George Baker, Robert Gibbons, and Kevin J. Murphy. Subjective performance measures in optimal incentive contracts. *Quarterly Journal of Economics*, 109(4):1125–1156, November 1994.

George Baker, Robert Gibbons, and Kevin J. Murphy. Relational contracts and the theory of the firm. *Quarterly Journal of Economics*, 117(1):39–84, February 2002.

Sandeep Baliga and Robert Evans. Renegotiation in repeated games with side-payments. *Games and Economic Behavior*, 33:159–176, 2000.

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.

Ken Binmore, Ariel Rubinstein, and Asher Wolinsky. The Nash bargaining solution in economic modelling. *RAND Journal of Economics*, 17(2):176–188, Summer 1986.

Adam Brandenburger and Harborne Stuart. Biform games. *Management Science*, 53(4):537–549, April 2007. URL http://mansci.journal.informs.org/cgi/content/abstract/53/4/537.

Lutz-Alexander Busch and Quan Wen. Perfect equilibria in a negotiation model. *Econometrica*, 63 (3):545–565, May 1995.

Pierre Cahuc, Fabien Postel-Vinay, and Jean-Marc Robin. Wage bargaining with on-the-job search: Theory and evidence. *Econometrica*, 74(2):323–364, March 2006. URL http://www.jstor.org/stable/3598804.

Wouter J. den Haan, Garey Ramey, and Joel Watson. Job destruction and propagation of shocks. *The American Economic Review*, 90(3):482–498, 2000. ISSN 00028282. URL http://www.jstor.org/stable/117339.

Peter A. Diamond. Wage determination and efficiency in search equilibrium. *The Review of Economic Studies*, 49(2):217–227, April 1982. URL http://www.jstor.org/stable/2297271.

Peter A. Diamond and Eric Maskin. An equilibrium analysis of search and breach of contract, i: Steady states. *The Bell Journal of Economics*, 10(1):282–316, Spring 1979. URL http://www.jstor.org/stable/3003332.

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

Christopher J. Flinn. Minimum wage effects on labor market outcomes under search, matching, and endogenous contact rates. *Econometrica*, 74(4):1013–1062, July 2006. URL http://www.jstor.org/stable/3805915.

Yuk-fai Fong and Jay Surti. The optimal degree of cooperation in the repeated prisoners' dilemma with side payments. *Games and Economic Behavior*, 67(1):277–291, September 2009.

James W. Friedman. A non-cooperative equilibrium for supergames. *Review of Economic Studies*, 38(1):1–12, January 1971.

Drew Fudenberg and Eric S. Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54(3):533–554, May 1986.

Drew Fudenberg, David K. Levine, and Eric S. Maskin. The folk theorem with imperfect public information. *Econometrica*, 62(5):997–1039, September 1994.

Sanford Grossman and Oliver Hart. The costs and benefits of ownership: A theory of lateral and vertical integration. *Journal of Political Economy*, 94:691–719, 1986.

Oliver Hart and John Moore. Property rights and the nature of the firm. *Journal of Political Economy*, 98:1119–1158, 1990.

Oliver D. Hart. *Firms, Contracts, and Financial Structure*. Clarendon, Oxford, 1995.

Oliver D. Hart and John Moore. A theory of debt based on the inalienability of human capital. *Quarterly Journal of Economics*, 109(4):841–879, November 1994.

Arthur J. Hosios. On the efficiency of matching and related models of search and unemployment. *The Review of Economic Studies*, 57(2):279–298, April 1990. URL http://www.jstor.org/stable/2297382.

Sebastian Kranz and Susanne Ohlendorf. Renegotiation-proof relational contracts with side payments. Working paper, October 2009.

Ehud Lehrer and Ady Pauzner. Repeated games with differential time preferences. *Econometrica*, 67(2):393–412, March 1999.

Jonathan Levin. Relational incentive contracts. *American Economic Review*, 93(3):835–857, June 2003.

W. Bentley MacLeod. Optimal contracting with subjective evaluation. *American Economic Review*, 92(1):216–240, March 2003.

W. Bentley MacLeod and James M. Malcomson. Implicit contracts, incentive compatibility, and involuntary unemployment. *Econometrica*, 57(2):447–80, March 1989.

Dale T. Mortensen and Christopher A. Pissarides. Job creation and job destruction in the theory of unemployment. *The Review of Economic Studies*, 61(3):397–415, 1994. ISSN 00346527. URL http://www.jstor.org/stable/2297896.

Giuseppe Moscarini. Job matching and the wage distribution. *Econometrica*, 73(2):481–516, March 2005. URL http://www.jstor.org/stable/3598795.

Roger B. Myerson. Values of games in partition function form. *International Journal of Game Theory*, 6(1):23–31, 1977.

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

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

Christopher A. Pissarides. Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *The American Economic Review*, 75(4):676–690, 1985. ISSN 00028282. URL http://www.jstor.org/stable/1821347.

Christopher A. Pissarides. Search, wage bargains and cycles. *The Review of Economic Studies*, 54 (3):473–483, July 1987. URL http://www.jstor.org/stable/2297570.

Roy Radner. Repeated principal-agent games with discounting. *Econometrica*, 53(5):1173–1198, September 1985.

Raghuram G. Rajan and Luigi Zingales. Power in a theory of the firm. *Quarterly Journal of Economics*, 113(2):387–432, May 1998.

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

Klaus M. Schmidt and Monika Schnitzer. The interaction of explicit and implicit contracts. *Economics Letters*, 48(2):193–199, May 1995.

Joel Watson. Contract and game theory: Basic concepts for settings with finite horizons. Working paper, July 2008. URL http://www.econ.ucsd.edu/publications/files/2006-01.pdf.

Asher Wolinsky. Matching, search, and bargaining. *Journal of Economic Theory*, 42(2):311–333, 1987.