A Theory of Disagreement in Repeated Games with Renegotiation*

David A. Miller† and Joel Watson‡
UCSD

Prepared for the NBER's Organizational Economics Working Group Meeting, November 20–21, 2009

Abstract

This paper analyzes repeated games with transferable utility, in which the players may negotiate cooperatively over their continuation strategies at the beginning of each period. In contrast to the renegotiation proofness literature, the model gives an explicit account of whether the players have reached an agreement in a given period, so there are feasible paths of play along which the players disagree. On a disagreement path, play may be jointly suboptimal. In a *contractual equilibrium*, the players cooperatively negotiate to play the continuation game optimally, splitting the surplus (according to fixed bargaining weights) relative to what they would have played under disagreement. Contractual equilibrium outcomes also arise in a class of models with noncooperative bargaining, under several assumptions on the endogenous meaning of cheap talk message(refining the set of subgame perfect equilibria). Contractual equilibria exist for all discount factors, and all such equilibria attain the same aggregate utility. The paper provides necessary and sufficient conditions for patient players to attain efficiency, as well as simple sufficient conditions. The allocation of bargaining power can dramatically affectaggregate utility. The theory extends naturally to games with imperfect public monitoring.

JEL Classifications: C71 (Cooperative games) , C72 (Noncooperative games) , C73 (Stochastic & dynamic games) , C78 (Bargaining theory; Matching theory)

*We thank seminar participants at Columbia, Florida International, the International Food Policy Research Institute, Penn, Toronto, UCLA, UCSD, UC Davis, USC, USC Marshall School of Business, Western Ontario, 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.

†Address: Dept. of Economics, 9500 Gilman Dr. #0508, La Jolla, CA 92093–0508. Home page: http://dss.ucsd.edu/~d9miller. 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.

‡Address: Dept. of Economics, 9500 Gilman Dr. #0508, La Jolla, CA 92093–0508. Home page: http://dss.ucsd.edu/~jwatson. Watson thanks the National Science Foundation (SES-0095207) and the National Oceanic and Atmospheric Administration (Southwest Fisheries Science Center) for financial support.

Contents

1 Introduction

Many contractual relationships exist over a long period of time and involve repeated productive interaction. In the absence of effective external enforcement, the parties must rely on self-enforcement to achieve cooperation (a setting termed *relational contracting*). Furthermore, an agreement on how to coordinate their behavior may be revisited by the players later in their relationship, so relational contracts have to deal with the potential problem of renegotiation.

To analyze these issues, we develop a theory of contracting and renegotiation in repeated games. Our main message is that whether and how players can agree to sustain cooperation is 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. The negotiation phase also allows voluntary monetary transfers. When they negotiate, the players may disagree, in which case they may play suboptimally until the next time they renegotiate. On the equilibrium path, however, they always agree on continuation payoffs that are Pareto optimal among those attainable in equilibrium.

Our theory describes behavior in the negotiation phase according to the generalized Nash (1950) bargaining solution, where the players use transfers to split the available surplus relative to the disagreement outcome, according to fixed bargaining weights. We introduce the concept of *contractual equilibrium*, which is a specification of actions, agreement paths, and disagreement paths satisfying incentive constraints in every action phase and the bargaining solution in every negotiation phase.[1] Our approach has bite if there are agreements the players can reach that are not attainable under disagreement. This feature is present in our model because we assume that disagreement occurs when a player makes a deviant transfer, implying that no transfers will be made under disagreement.

We also specify 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 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 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.

Our theory achieves two goals: (i) providing a coherent account of agreement and disagree-

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

ment in repeated games, where bargaining power plays a central role; and (ii) providing a technical apparatus that is straightforward to apply. We prove that if the stage game is finite then a contractual equilibrium always exists, regardless of the discount factor,[2] and that the welfare level (joint payoff) 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 how the attainable welfare level depends critically on the players' relative bargaining weights.

The features that distinguish our theory from the previous literature are the explicit account of negotiation and the role of disagreement. The renegotiation proofness literature (initiated by Bernheim and Ray 1989, Farrell and Maskin 1989, and Pearce (1987); and further developed by Asheim 1991; Abreu, Pearce, and Stacchetti 1993; Abreu and Pearce (Unknown); Bergin and MacLeod 1993; Ray 1994; and others) allows the players, after each history, to jointly deviate from their planned continuation strategies to another continuation equilibrium in a particular class. A renegotiation-proof equilibrium is a subgame perfect strategy profile away from which the players would never jointly deviate. Baliga and Evans (2000), Levin (2003), Fong and Surti (2009), and Kranz and Ohlendorf (2009) particularly address renegotiation proofness in games with transferable utility, as do some works in the relational contracts literature (e.g., Levin 2003). Renegotiation-proofness does not address the possibility of disagreement, but in the context of our setting it can be viewed as requiring the players to play the same way under both agreement and disagreement. In such an equilibrium there is no surplus to negotiate over, so there is no role for the exercise of bargaining power.

In our view, it is reasonable to expect that the players might not play optimally when they are in disagreement. Empirically, it is common to see negotiating parties suffer from costly delay while trying to reach an agreement. Indeed, it is not uncommon to see parties actively burn surplus while in disagreement, such as when workers go on strike or their firms lock them out.[3] Conceptually, the literature on bargaining in one-shot settings assumes as a matter of course that outcomes under disagreement are suboptimal.

In this preliminary and incomplete draft, Section 2 describes repeated games with renegoti-

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

[3]To be sure, in a contractual equilibrium the players always agree on the equilibrium path, so disagreement should never be observed. We expect that the model would have to incorporate overoptimism, private information, or some other similar element to generate disagreement on the equilibrium path; such extensions are outside the scope of this paper.

ation and transfers. In Section 3 we define our equilibrium concept and illustrate aspects of the construction using a prisoners' dilemma example. Section 4 contains details on the existence and characterization of the contractual equilibrium value set. Section 5 discusses when efficient equilibria exist, and shows how to construct them. Section 6 shows that bargaining power can dramatically affect the level of welfare the players can attain. Section 7 extends the analysis to settings with imperfect public monitoring. Section 8 provides noncooperative foundations for contractual equilibrium. We plan to address the case of $n > 2$ players in a later draft. Some discussion, details, proofs, examples, and citations are missing from this draft and will be added later.

2 Repeated games with renegotiation and transfers

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 jointly determine their continuation strategies and can voluntarily make immediate monetary transfers. Money enters their payoffs quasi-linearly. In the action phase, the players simultaneously choose actions in the stage game. For most of the paper, we assume that these actions are commonly observed, but in Section 7 we consider games with imperfect monitoring.

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 selection of 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, loosely, two mappings: one from histories ending in realizations of the public randomization device to agreements, and one from histories ending with an agreement or a disagreement to mixed action profiles in the stage game. An agreement comprises a continuation payoff vector and a monetary transfer (which is implemented immediately).

A *contractual equilibrium* is a strategy profile in which, after every history, the players reach a Nash bargain in the negotiation phase, play sequentially rational actions in the action phase, and attain 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 Contractual equilibrium

We next define contractual equilibrium recursively, and provide some basic characterizations. We start with some notation. Let $\mathbb{R}^2_0 \equiv \{m \in \mathbb{R}^2 \mid m_1 + m_2 = 0\}$ be the set of budget-balanced transfer vectors. For any product set H, let $\Delta^{\mathrm{U}} H$ denote the set of 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. Let $Y \subset \mathbb{R}^2$ denote a set of feasible continuation values from the negotiation phase in a given period. In essence, Y is the set of continuation values over which the players negotiate. Let $\underline{Y}$ be the subset of Y that can be achieved without making an immediate monetary transfer. We construct V, Y, and $\underline{Y}$ recursively below.

The dynamic programming equation characterizing an agreement is

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

where $y \in Y$ 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

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

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 $\langle A, (1 - \delta)u(\cdot) + \delta g(\cdot) \rangle$.

Given some set V of continuation values available in the following period, the set of payoffs that can be agreed upon in the current negotiation phase is given by the operator C:

$$C(V) \equiv \{y \in \mathbb{R}^2 \mid \exists\, m \in \mathbb{R}^2_0, g : A \to V, \text{ and } \alpha \in \Delta^{\mathrm{U}} A \text{ s.t. } h \text{ enforces } \alpha \text{ and (1) holds}\}. \tag{3}$$

Under disagreement, there are no transfers, so the set of payoffs that can arise following a disagreement is given by the operator $\underline{C}$:

$$\underline{C}(V) \equiv \{y \in \mathbb{R}^2 \mid \exists\, g : A \to V \text{ and } \alpha \in \Delta^{\mathrm{U}} A \text{ s.t. } g \text{ enforces } \alpha \text{ and (2) holds}\}. \tag{4}$$

Note that both C and $\underline{C}$ map subsets of $\mathbb{R}^2$ to subsets of $\mathbb{R}^2$.

If V describes the set of achievable continuation payoffs from the next period, then $Y = C(V)$ is the implied set of achievable continuation values from the negotiation phase in the current period and $\underline{Y} = \underline{C}(V)$ are the values that can be achieved without an immediate transfer. Note that $C(V)$ is closed under constant transfers, meaning that $y \in C(V)$ and $m \in \mathbb{R}^2_0$ imply that $y + m \in C(V)$. Also, if V is compact then it must be that $\max\{y_1 + y_2 \mid y \in C(V)\}$ exists. Furthermore, since the Nash equilibrium graph is closed, $\underline{C}$ preserves compactness.

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

$$y_1 + y_2 = \max_{y' \in Y} y'_1 + y'_2 \quad \text{and} \quad \pi_2(y_1 - \underline{y}_1) = \pi_1(y_2 - \underline{y}_2). \tag{5}$$

The second part of this condition is equivalent to

$$y_i = \underline{y}_i + \pi_i(y_1 + y_2 - \underline{y}_1 - \underline{y}_2) \text{ for each player } i. \tag{6}$$

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:

$$B(Y, \underline{Y}) \equiv \{y \in Y \mid \exists\, \underline{y} \in \underline{Y} \text{ s.t. (5) holds}\}. \tag{7}$$

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 is also compact.

We next put the three conditions together to describe the formal relation between V, Y, and $\underline{Y}$. 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 $Y = C(V)$, $\underline{Y} = \underline{C}(V)$, and including the public randomization device, $V = \operatorname{co} B(Y, \underline{Y})$, where "co" denotes "convex hull." 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 $V = \operatorname{co} B(C(V), \underline{C}(V))$.

When discussing an NSG set V, we will refer to $Y = C(V)$ and $\underline{Y} = \underline{C}(V)$ without reference to C or $\underline{C}$. The next lemma follows from the properties of B, C, and $\underline{C}$ 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 our analysis, we will focus on the case in which V contains its endpoints and is therefore closed; we will later handle the case in which V does not contain its endpoints. 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 elements of $B(Y, \underline{Y})$. That is, their presence in $V = \operatorname{co} B(Y, \underline{Y})$ is not only as convex combinations of other points in $B(Y, \underline{Y})$ because this would imply that z^1 and z^2 are not endpoints. 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 (5).

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 $v_2 \geq \hat{v}_2$.

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

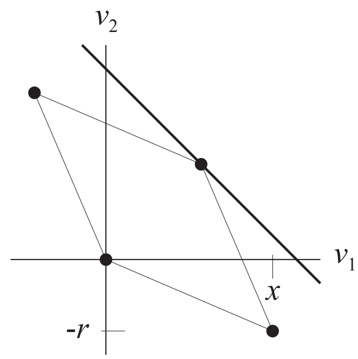

Figure 1: The Prisoners' Dilemma with Transfers. The stage game parameters satisfy $x > r > 0$. The Pareto frontier is a line with slope -1 through the point $(1, 1)$.

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 $z^1_2 \geq \hat{z}^1_2$ and $z^2_1 \geq \hat{z}^2_1$. Our main definition combines negotiated self-generation with dominance to form a stringent 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. Loosely speaking, a contractual equilibrium strategy profile is a selection $(y, \underline{y}) \in C(V^*) \times \underline{C}(V^*)$ for each $v \in V^*$, a selection $\alpha \in \Delta^{\mathrm{U}} A$ for each $y \in C(V^*) \cup \underline{C}(V^*)$, and a selection $m \in \mathbb{R}^2_0$ for each $\underline{y} \in \underline{C}(V^*)$ that are consistent with the forgoing construction. Note that since we have defined a hybrid cooperative-noncooperative game, the selections of y and $\underline{y}$ represent joint actions, while the selections of m and α represent vectors of individual actions.

3.1 Example: Prisoners' Dilemma

For an example, consider the repeated prisoners' dilemma with transfers and renegotiation. Suppose that the players have equal bargaining weights. The stage game and feasible payoff set are pictured in Figure 1. We will later identify conditions under which z^2 is different than z^1 (that is, V^* has strictly positive length). The conditions relate to whether it is possible to support play of

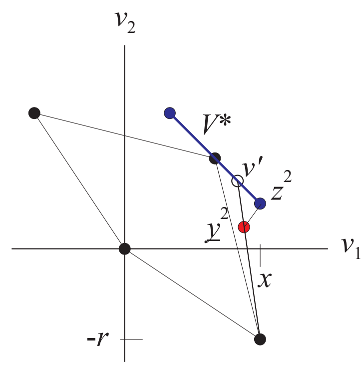

FIGURE 2: PRISONERS' DILEMMA z^2 CONSTRUCTION. The endpoint z^2 is attained by playing CC in the stage game and using a transfer to split the surplus relative to the disagreement point $\underline{y}^2$. The disagreement point $\underline{y}^2$, in turn, is attained by playing DC in the stage game and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

(C, D) and (D, C) in a given period. For now, suppose that it is possible to support these action profiles.

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 perpendicular to the vector π — in other words, in the direction $(\pi_2, -\pi_1)$. This disagreement point, in turn, will be the weighted average of $u(\mathrm{D}, \mathrm{C}) = (x, -r)$ 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, because this best favors player 1. However, importantly v' cannot equal z^2 because we need room to punish player 2 if he were to deviate from the action profile (D, C). Thus, we specify that the players select (D, C) 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.

4 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, which is* *defined by $\langle A, u, \delta, \pi\rangle$, and assume A is finite and $\delta \in [0, 1)$. This game has a unique CEV set V^*.*

The rest of this section contains the proof of the theorem, which also shows how to calculate 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 $z_1^1 + z_2^1$ and $z_1^2 - z_1^1$. We will show the steps for z^2; the logic is the same for z^1.

Given $V = \mathrm{co}\{z^1, z^2\}$, consider the following problem, which identifies player 1's most preferred outcome of negotiation utilizing continuation values from the V associated with the next period:

$$\max \left\{ v_1 \middle| v \in B\big(C(V), \underline{C}(V)\big) \right\} \tag{8}$$

Because elements in $B(C(V), \underline{C}(V))$ correspond to various disagreement points in $\underline{C}(V)$, this maximization problem can be expressed using the bargaining solution (Expression 5) and in terms of the disagreement point as the choice variable. Letting

$$L \equiv \max_{y \in C(V)} \; y_1 + y_2, \tag{9}$$

note that the bargaining solution gives player 1 the value $\underline{y}_1 + \pi_1(L - \underline{y}_1 - \underline{y}_2)$, which equals $\pi_2 \underline{y}_1 - \pi_1 \underline{y}_2 + \pi_1 L$. The maximization problem can therefore be written as:

$$\begin{aligned} &\max_{\underline{y}, h, \alpha} \; \pi_2 \underline{y}_1 - \pi_1 \underline{y}_2 + \pi_1 L, \\ &\text{s.t.} \; \begin{cases} \underline{y} = (1-\delta)u(\alpha) + \delta h(\alpha), \\ h \text{ and } \alpha \text{ satisfy Condition CA.} \end{cases} \end{aligned} \tag{10}$$

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 $g_1(a) + g_2(a) - z_1^2 - z_2^2 = 0$. Thus, we have $g_1(a) = \eta(a) + z_1^2$ and $g_2(a) = z_2^2 - \eta(a)$. Also, the constraint that $g(a) \in \mathrm{co}\{z^1, z^2\}$ is equivalent to the requirement that $\eta(a) \in [z_1^1 - z_1^2, 0]$.

Using these facts to substitute for g, the maximization problem is written as

$$\max_{\eta,\alpha} \quad (1-\delta)\big(\pi_2 u_1(\alpha) - \pi_1 u_2(\alpha)\big) + \delta(\pi_2 z_1^2 - \pi_1 z_2^2)\big),$$
$$\text{s.t.} \begin{cases} \eta : A \to [z_1^1 - z_1^2, 0] \text{ extended to } \Delta^{\mathrm{U}} A, \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \big\langle A, (1-\delta)u(\cdot) + \delta(\eta(\cdot), -\eta(\cdot)) + z^2 \big\rangle. \end{cases} \tag{11}$$

Remember that this is equivalent to the maximization problem in (8). Because V is NSG, we know that the value at the solution to this problem is precisely z_1^2. For the same reason, we know that $z_1^2 + z_2^2 = L$, which we can use to substitute for z_2^2 in the second bracketed term of the objective function. Combining some of the z_1^2 terms, using the fact that $\pi_1 + \pi_2 = 1$, and dividing by $(1-\delta)$, we obtain:

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

This expression can be simplified a bit further. In the Nash equilibrium condition for α (the implied game in the constraint), we can remove the constant z^2 and divide by $(1-\delta)$; clearly, this does not change the equilibrium incentive conditions. We also change variables by writing $\varphi = \frac{\delta}{1-\delta}\eta$. This yields the following condition that partly characterizes z^2:

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

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

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

Note that $\overline{\gamma}$ is a function of the *payoff span*, $d \equiv z_1^2 - z_1^1$, which is the vertical distance between z^1 and z^2 (the difference between each player's most and least preferred points).

By similar calculations, we obtain a partial characterization of z^1:

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

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

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

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

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

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

In other words, this lemma says that if the length (payoff span) of $\tilde{V}$ strictly exceeds that of $\hat{V}$ then $\tilde{V}$ dominates $\hat{V}$.

Proof. Suppose that $\tilde{z}_1^2 - \tilde{z}_1^1 > \hat{z}_1^2 - \hat{z}_1^1$. The larger length of $\tilde{V}$ can support weakly more mixed actions in the stage game as equilibria than can $\hat{V}$ simply because $\tilde{V}$ 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 only amount to constants in the players' payoffs), only their relative lengths. Thus, any mixed action that can be supported in the context of $\hat{V}$ can also be supported in the context of $\tilde{V}$. One can see that, for any NSG set, the level $z_1^2 + z_2^2$ is the highest welfare level that can be supported in the stage game. This implies that $\tilde{z}_1^2 + \tilde{z}_2^2 \geq \hat{z}_1^2 + \hat{z}_2^2$; that is, the welfare level of $\tilde{V}$ weakly exceeds the welfare level of $\hat{V}$.

Furthermore, the function $\bar{\gamma}$ is clearly increasing. Equation 13 and the larger length of $\tilde{V}$ therefore imply that $\tilde{z}_1^2 \geq \hat{z}_1^2$. Also note that, using the facts $\pi_1 + \pi_2 = 1$ and $z_1^1 + z_2^1 = z_1^2 + z_2^2$, we can rearrange Equation 15 to form:

$$z_2^1 = \pi_2(z_1^2 + z_2^2) - \underline{\gamma}(z_1^2 - z_1^1). \tag{17}$$

Observe as well that $\underline{\gamma}$ is decreasing. That $\tilde{V}$ has a larger length than $\hat{V}$, and since $\tilde{z}_1^2 + \tilde{z}_2^2 \geq \hat{z}_1^2 + \hat{z}_2^2$, we thus know that $\tilde{z}_2^1 \geq \tilde{z}_2^1$, which completes the proof that $\tilde{V}$ dominates $\hat{V}$.

The case of $\tilde{z}_1^2 - \tilde{z}_1^1 = \hat{z}_1^2 - \hat{z}_1^1$ is easily handled as well. Here, the two sets must have the same welfare level and, because their payoff spans are equal, they also must have the same endpoints and are therefore equivalent. $\square$

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

$$\Gamma(d) \equiv \bar{\gamma}(d) - \underline{\gamma}(d) \tag{18}$$

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.

Note 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 $\bar{\gamma}$ 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:

$$\begin{aligned} L^* \equiv \max_{\varphi,\alpha} \quad & u_1(\alpha) + u_2(\alpha), \\ \text{s.t.} \quad & \begin{cases} \varphi : A \to [0, \frac{\delta}{1-\delta} d^*], \text{ extended to } \Delta^{\mathrm{U}} A, \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \langle A, u(\cdot) + (\varphi(\cdot), -\varphi(\cdot)) \rangle. \end{cases} \end{aligned} \tag{19}$$

Then we obtain z_1^{2*} and z_1^{1*} using Equations 13 and 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. This is because new Nash equilibria could emerge in the stage game. 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.

5 Efficiency

Inspection of $\overline{\gamma}$ and $\underline{\gamma}$ reveals that Γ is bounded and increasing in δ. Thus, the payoff span d^* of the CEV set is increasing (weakly) in the discount factor. Furthermore, if $\Gamma(\infty) > 0$ then we have $d^* > 0$ for δ close enough to one. This, in turn, means that any action profile can be supported in a single period when players are patient. This proves 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 $V^* = \{u(\alpha^*)\}$, where α^* is the joint-payoff-maximizing Nash equilibrium of the stage game.*

The result is illustrated by the Prisoners' Dilemma case (recall Figure 1). Examining the maximization and minimization problems that define $\overline{\gamma}$ and $\underline{\gamma}$, one can easily show that $\overline{\gamma}(\infty) > 0$ and $\underline{\gamma}(\infty) < 0$ if and only if $x > r$. In this case, the contractual equilibrium is efficient for sufficiently large δ. Otherwise, $\overline{\gamma}(\infty) = 0$ and $\underline{\gamma}(\infty) = 0$ and we have $V^* = \{(0,0)\}$ regardless of δ. The condition on x and r is apparent by checking how to support play of (D, C) in the case of $\overline{\gamma}$ and (C, D) in the case of $\underline{\gamma}$.

Next we provide methods for constructing explicit contractual equilibria 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 machine strategy profile, 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, 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, stay in state ω^i. If player j deviates unilaterally, go to state ω^j.*

Proof. 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 $\hat{d} = \hat{z}^2_1 - \hat{z}^1_1$ 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^*.

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 $du_i(\alpha)/d\alpha_{-i} < -1/k$.*[4] *Then there exists an efficient contractual equilibrium if the players are sufficiently patient.*

Proof. 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 $\epsilon > 0$ let $\alpha_i^{-i}(\epsilon) \equiv \alpha_i^* + \epsilon$ and $\alpha_i^i(\epsilon) \equiv BR_i(\alpha_{-i}^i(\epsilon))$. Then for all $\epsilon < \eta/2k$, $\alpha_i^* - \epsilon k < \alpha_i^i(\epsilon) < \alpha_i^* + \epsilon k$.

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 $\left|u_i\left(\alpha_i^i(\epsilon), \alpha_{-i}\right) - u_i\left(\alpha_i^{-i}(\epsilon), \alpha_{-i}\right)\right|$ 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 $u_i\left(\alpha_i, \alpha_{-i}^{-i}(\epsilon)\right) - u_i\left(\alpha_i, \alpha_{-i}^i(\epsilon)\right) > 0$ is on the order of at least ϵ. Hence, for $\epsilon > 0$ 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. □

As an example, the standard Cournot duopoly game does not satisfy the conditions of Corollary 1, since both firms earn zero profits whenever one firm is minimaxed. However, it satisfies the conditions of Theorem 4, and therefore has an efficient contractual equilibrium if the firms are sufficiently patient.

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

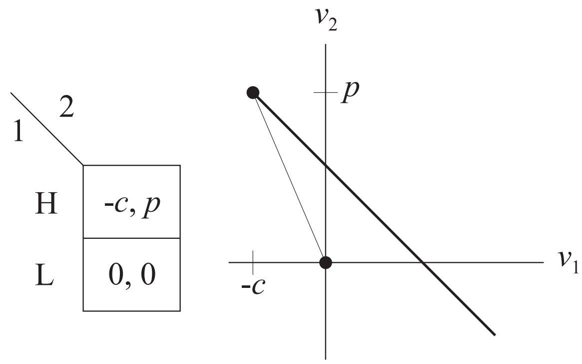

Figure 3: The Principal-Agent game. Output p is received by player 2 when player 1 selects effort H; otherwise output is 0. The cost of effort H is $c \in (0, p)$.

6 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 this section we show by means of a simple example that the allocation of bargaining power affects whether and when efficiency is attainable. Then we conjecture (and outline an approach to proving) that in symmetric games the highest welfare level is attained when bargaining power is extremely unequal.

6.1 Example: The Principal-Agent game

Consider a repeated principal-agent game with no external enforcement. The stage game is as shown in Figure 3. In this game, the agent (player 1) chooses whether to exert high or low effort in each period; high effort entails a personal cost of c. The principal receives the output created by the agent's action. If the agent selects low effort, then output is zero. If the agent selects high effort then output is p. Assume that $p > c$.

The construction of V^* is straightforward. It is easy to calculate that, for all π, the point z^1 gives player 1 exactly zero and, for large enough d, this is accomplished by having player 1 select H in the disagreement point. On the other hand, player 1 selects L in the disagreement point that is associated with z^2. That is, player 1 is punished using a disagreement point in which he would exert high effort and he is rewarded using a disagreement point in which he selects L. Figure 4 illustrates the construction.

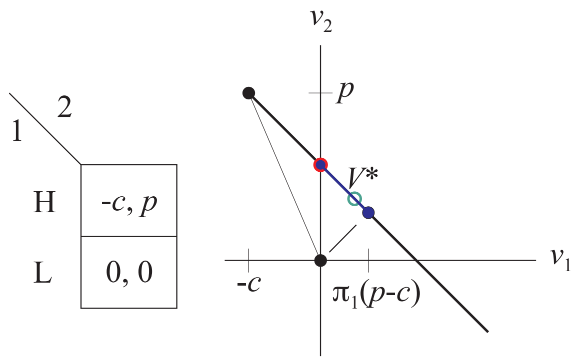

Figure 4: Construction of z^1 and z^2 for the Principal-Agent game. The point z^1 is an agreement formed relative to itself as its own disagreement point, so there is no surplus to split. It is attained by playing H and continuing with the green circle, while switching to z^1 if L is chosen instead. The point z^2 is the agreement formed relative to the disagreement point $(0, 0)$. It is attained by playing H and continuing with z^2, while switching to an interior point (not shown) in V^* if L is chosen instead.

		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 5: The Prisoners' Double Dilemma game.

In this game, we have $\bar{\gamma}(\infty) = 0$ (where L is chosen) and $\underline{\gamma}(\infty) = -\pi_1(p - c)$, so $\Gamma(\infty) = \pi_1(p - c)$. Efficiency (play of H each period) is attainable by patient players if and only if $\pi_1 > 0$. Moreover, the threshold discount factor for attaining efficiency decreases in π_1, because it affects z^2 but not z^1. This is related to the fact that the disagreement point from which the players negotiate to z^2 is inefficient, whereas the z^1 is itself the disagreement point at the other side of V^* (so there is no bargaining surplus there).

6.2 Symmetric games

Even in symmetric games, bargaining power can have a dramatic effect on the attainable welfare level. An instructive example is the Prisoners' Double Dilemma shown in Figure 5. 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 $\pi = \left(\frac{1}{2}, \frac{1}{2}\right)$, the two stage game equilibria DD and EE 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 DD and EE has a strictly positive payoff span, and thus efficiency is attainable for sufficiently high δ.

Building on this example, we conjecture that in symmetric games the welfare level is maximized under maximally unequal bargaining power.

Conjecture 1. *In a symmetric game, L^* is maximized at both $\pi = (0, 1)$ and $\pi = (1, 0)$.*

Proof idea. For $\pi = \left(\frac{1}{2}, \frac{1}{2}\right)$, disagreement points $\underline{y}^1$ and $\underline{y}^2$ that support V^* are w.l.o.g. symmetric. Geometrically show that using these disagreement points yields the same span d regardless of π.

For $\pi_1 > \frac{1}{2} > \pi_2$, the optimal disagreement points may be asymmetric, which means the span they support may vary with π. But for any two optimal disagreement points for some $\pi_1 > \frac{1}{2} > \pi_2$, there are optimal mirror-image disagreement points for the reciprocal bargaining weights, which are enforceable on the same span. Geometrically show that this implies the span enforced using the best two of these four disagreement points is weakly increasing in $\pi_i \geq \frac{1}{2}$.

Collectively these facts imply that the welfare level in contractual equilibrium is weakly increasing in $\pi_i \geq \frac{1}{2}$. A complete proof awaits the next draft of this paper. □

7 Imperfect public monitoring

In this section, we describe a more general model that allows for imperfect public monitoring. A game in this class is defined by:

- A stage game, featuring a finite set of action profiles $A = A_1 \times A_2$, a finite set of public signals S, a signal distribution function $p : S \times A \to [0, 1]$, and payoff functions $u_1 : A_1 \times S \to \mathbb{R}^2$ and $u_2 : A_2 \times S \to \mathbb{R}^2$;
- A discount factor $\delta \in [0, 1)$;
- Bargaining weights $\pi = (\pi_1, \pi_2)$, with $\pi_1 + \pi_2 = 1$ and $\pi_1, \pi_2 \geq 0$.

In the stage game, player 1 selects an action $a_1 \in A_1$ and simultaneously player 2 selects an action from $a_2 \in A_2$, so the action profile is $a = (a_1, a_2)$. Public signal s is then realized with probability $p(s, a)$. The players both observe s but they do not observe each other's actions. We write $u(a, s) = u_1(a_1, s), u_2(a_2, s))$. We extend p to the space of mixed actions, so, letting $\alpha(a)$ denote the probability that mixed action profile α assigns to action profile a, we have $p(s, \alpha) = \sum_{a \in A} \alpha(a) p(s, a)$.

This is essentially a game of observed actions (the class described earlier in this paper) if $S = A$ and if, for each action profile a, we have $s = a$ with probability one. The stage game has imperfect public monitoring if a player cannot determine the exact action played by the other player given her own action and the public signal.

The other aspects of the model are exactly as in the case with observed 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 conditioned on the public signals rather than 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."

Most components of the analysis are unchanged, including the definition and meaning of V, Y, and $\underline{Y}$. The construction of a continuation value y from the negotiation phase (Equation 1 in the basic model) is now:

$$y = m + \sum_s p(s, \alpha)[(1 - \delta)u(\alpha, s) + \delta g(s)], \tag{20}$$

where $m \in \mathbb{R}^2_0$ is the immediate transfer, $\alpha \in \Delta^{\mathrm{U}} A$ is a mixed action profile to be played in the current period, and g gives the continuation value from the start of the next period as a function of the public signal in the current period.

Condition CA becomes:

Condition CA′: The function g maps S to V. Further, α is a Nash equilibrium of the game with action-profile space A and where, for each $a \in A$, the payoffs are given by $\sum_{s \in S} p(s, a)[(1 - \delta)u(a, s) + \delta g(s)]$.

Operators C and $\underline{C}$ are revised as follows:

$$C(V) \equiv \{y \mid \text{There exist } m \in \mathbb{R}^2_0,\ g, \text{ and } \alpha \in \Delta^{\mathrm{U}} A \text{ such that Condition CA}' \text{ and Equation 20 hold}\}$$

and

$$\underline{C}(V) \equiv \{y \mid \text{There exist } g \text{ and } \alpha \in \Delta^{\mathrm{U}} A \text{ such that Condition CA}' \text{ holds and Equation 20 is satisfied with } m = (0,0)\}.$$

The function B requires no modification for the general model. As before, V is said to satisfy negotiated self-generation (NSG) if $V = \operatorname{co} B(C(V), \underline{C}(V))$. Lemma 1 is still valid and the definition of dominance and contractual equilibrium are the same as before.

With A and S finite, we can guarantee existence. The analysis proceeds just as for the case with observed actions worked out in Section 4. Functions $\bar{\gamma}$ and $\underline{\gamma}$ are defined as follows:

$$\begin{aligned}\bar{\gamma}(d) \equiv \max_{\varphi,\alpha} \quad & \sum_{s\in S} p(s,\alpha)\left[\pi_2 u_1(\alpha_1,s) - \pi_1 u_2(\alpha_2,s) + \varphi(s)\right], \\ \text{s.t.} \quad & \begin{cases} \varphi : S \to \left[-\frac{\delta}{1-\delta}d, 0\right], \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \left\langle A, \sum_{s\in S} p(s,\cdot)\left[u(\cdot,s) + (\varphi(s), -\varphi(s))\right]\right\rangle \end{cases}\end{aligned} \tag{21}$$

and

$$\begin{aligned}\underline{\gamma}(d) \equiv \min_{\varphi,\alpha} \quad & \sum_{s\in S} p(s,\alpha)\left[\pi_2 u_1(\alpha_1,s) - \pi_1 u_2(\alpha_2,s) + \varphi(s)\right], \\ \text{s.t.} \quad & \begin{cases} \varphi : S \to \left[0, \frac{\delta}{1-\delta}d\right], \\ \alpha \in \Delta^{\mathrm{U}} A \text{ is a Nash equilibrium of } \left\langle A, \sum_{s\in S} p(s,\cdot)\left[u(\cdot,s) + (\varphi(s), -\varphi(s))\right]\right\rangle. \end{cases}\end{aligned} \tag{22}$$

As before, the contractual equilibrium value set V^* is determined by finding the maximal fixed point of $\Gamma = \bar{\gamma} - \underline{\gamma}$.

7.1 Welfare levels for patient players

With regard to efficiency, there is an extra condition relative to the case of observed actions: that the public signal distribution p is sufficient to detect deviators relative to the desired stage-game action profile. We have:

Theorem 5. *Consider a repeated game with imperfect public monitoring, renegotiation, and transfers. If* $\bar{\gamma}(\infty) > \underline{\gamma}(\infty)$ *then* $d^* > 0$ *and the level of* V^* *is given by:*

$$\max_{\alpha} \sum_{s\in S} p(s,\alpha)[u_1(\alpha_1,s) + u_2(\alpha_2,s)], \tag{23}$$

subject to α *being enforceable with respect to continuation values on the line* $\mathbb{R}^2_0$ *(that is,* α *is "orthogonally enforceable" in direction* $(1,1)$*).*

Proof. This follows from the same argument used to prove Theorem 2. □

Thus, two conditions are required to for efficient contractual equilibria: We need the payoff span of V^* to be strictly positive and we need the efficient action profile to be enforceable using continuation values on a line of slope -1.

7.2 Example: The Imperfect Principal-Agent game

To illustrate, we augment the Principal-Agent game with imperfect information. Suppose that when the agent chooses high effort, the principal receives output 1 with probability $p < 1$, and receives zero otherwise. The principal does not observe the agent's effort. Then it is simple to construct a contractual equilibrium similar to that of Section 6.1, except that high output, rather than high effort, should lead to a high continuation value for the agent, and similarly for low output.

8 Noncooperative foundations

In this section, we provide noncooperative foundations for contractual equilibrium. We start by specifying a fully noncooperative model of interaction in the repeated game (including bargaining in each period). For simplicity, we focus on the setting of perfect monitoring; the extension to imperfect public monitoring is straightforward. Suppose that bargaining takes place by way of a random-proposer, ultimatum-offer protocol. Each period comprises four phases: (1) the public randomization phase, (2) the bargaining phase, (3) the voluntary transfer phase, and (4) the (stage game) action phase.

In the action phase, the players play the stage game $\langle A, u \rangle$ defined in Section 2. 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 sum transfer is denoted m. 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." Then the offerer selects a statement from some language set Λ. The other player (the "responder") then says "yes" or "no."[5]

[5]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 the following alternating-offer protocol in "stop time" (where arbitrarily many rounds of offers and responses happen in infinitesimal time compared to the length of a period): In each round of bargaining, one player is randomly selected to make a proposal and the other player then responds with "yes" or "no." If the responder says "yes", then the bargaining phase ends and they proceed to the transfer phase. If the responder says "no", then with some fixed positive probability the bargaining process exogenously breaks down and they proceed to the transfer phase; otherwise bargaining continues to another round.

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 is to assume that Λ contains descriptions of all continuation strategy profiles in the game. Since such a construction is circular (strategies specify history-dependent statements in the negotiation phase, and these statements include the description of a strategy profile), it leads to an interesting technical issue regarding the existence of a "universal language." 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 the players as a suggested continuation value.

We shall take the simpler approach and assume that

$$\Lambda \equiv \{\text{"discontinue"}\} \cup \mathbb{R}^2 \times \mathbb{R}_0^2. \tag{24}$$

At this point in the development of the model, statements in Λ have no particular meaning, but we make assumptions below that imply meaning in equilibrium. Here is the flavor of what follows: We shall regard the statement "discontinue" as expressing that the offerer will not continue with the current agreement path. Furthermore, a statement of some point $\lambda = (w, m) \in \mathbb{R}^2 \times \mathbb{R}_0^2$ will be interpreted as a specific suggestion about how the players should coordinate in the continuation of the game. By stating λ, 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 $y = m + w$ from the negotiation phase.

Let S be the set of all subgame-perfect equilibria of our fully noncooperative model, in behavioral strategies (to allow for mixing). 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, 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 $\lambda = (w, m)$

and "yes" as indicating that the players have accepted the proposal $\lambda = (w, m)$ 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 step involves building notions of meaningful disagreement and agreement. 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

$$H^t \equiv (\Omega \times \{1, 2\} \times \Lambda \times \{\text{"yes"}, \text{"no"}\} \times \mathbb{R}_0^2 \times A)^t. \tag{25}$$

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

$$\hat{H} \equiv H \times \Omega \times \{1, 2\} \times \Lambda \times \mathbb{R}_0^2. \tag{26}$$

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

First, we 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 prior to the current offerer selection* if $\hat{h} = (h; \omega, k, \lambda, m)$ and $\hat{h}' = (h; \omega, k', \lambda', m')$ 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 prior 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 called the *disagreement point.*

Let S^{D} be the subset of S^{C} that satisfy no-fault disagreement. Since disagreement events occur only off the equilibrium path (by definition of S^{C}), the continuation values supported by S and S^{D} are the same, so the notion of no-fault disagreement does not on its own imply a refinement of supported payoffs. However, no-fault disagreement in combination with our next condition will result in a refinement.

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 $\hat{h} = (h; \omega, k, \lambda, m) \in \hat{H}$. Then $\hat{h}$ is included in $\hat{H}^{\mathrm{I}}(s)$ if and only if:

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

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

(iii) 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 agree 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 $\tilde{L}$ be the supremum of levels (joint continuation values) over all strategies in S' and all histories in H. That is, letting $\tilde{v}(s, h)$ denote the continuation value under strategy s following history h,

$$\tilde{L} \equiv \sup\{\tilde{v}_1(s, h) + \tilde{v}_2(s, h) \mid s \in S',\ h \in H\}. \tag{27}$$

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

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 $\hat{h} = (h; \omega, k, \lambda, m) \in \hat{H}$. Then $\hat{h}$ is included in $\hat{H}^{\mathrm{E}}(s, S')$ if and only if:

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

(ii) $m = m'$, and

(iii) 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 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.[6] Our next result shows that payoffs attained by equilibria in S^* are the same as those attained by a contractual equilibrium.

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.*

Proof. 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^* must attain the same welfare level as V^*.

[6]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.

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 $\underline{y}^1 \in \underline{C}(V^*)$ such that

$$z^1 = \underline{y}^1 + \pi(L^* - \underline{y}_1^1 - \underline{y}_2^1). \tag{28}$$

From the definition of $\underline{C}$ (recall Expression 4), 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

$$\underline{y}^1 = (1 - \delta) u(\alpha^1) + \delta g^1(\alpha^1). \tag{29}$$

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 $w \in \underline{C}(V^*)$ 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 $w \in \underline{C}(V^*)$ 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 $w \notin \underline{C}(V^*)$ 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 $\underline{C}(V^*)$, 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. For any set $\underline{Y} \subset \mathbb{R}^2$ and any level $L \in \mathbb{R}$, define:

$$\tilde{B}(L, \underline{Y}) \equiv \{\underline{y} + \pi(L - \underline{y}_1 - \underline{y}_2) \mid \underline{y} \in \underline{Y} \text{ satisfying } \underline{y}_1 + \underline{y}_2 \leq L\}. \tag{30}$$

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 $\underline{C}(V)$, we have $V \subset B(L, \underline{C}(V))$. 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 $\underline{y} \in \underline{C}(V)$. 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^* and so $L \leq L^*$. Furthermore, it is the case that 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 $\underline{C}(V^*)$. 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. □

References

Dilip Abreu and David Pearce. A perspective on renegotiation in repeated games. Working paper, Unknown.

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

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

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

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.

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

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.

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

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

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.

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

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.