A Theory of Disagreement in Repeated Games with Bargaining

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

October 13, 2011

Abstract

This paper proposes a new approach to the problem of equilibrium selection in repeated games with transfers, by supposing that in each period the players bargain over how to play. Although the bargaining phase is cheap talk (which follows a generalized alternating-offer protocol), sharp predictions arise from three axioms. Two axioms allow the players to meaningfully discuss whether to deviate from their plan; the third embodies a "theory of disagreement"—that play under disagreement should not vary with the manner in which bargaining broke down. Equilibria satisfying these axioms exist for all discount factors and are simple to construct, and all equilibria attain the same joint value. Optimal play under agreement generally requires suboptimal play under disagreement. Whether patient players attain efficiency depends on both the stage game and the bargaining power that they derive from the details of the bargaining protocol. The theory extends naturally to games with imperfect public monitoring and heterogeneous discount factors, and yields new insights into classic relational contracting questions.

JEL Classifications: C71 , C72 , C73 , C78

*We thank Sylvain Chassang, Marina Halac, Jin Li, Larry Samuelson, and Garey Ramey for valuable comments and suggestions, along with seminar participants at Carlo Alberto, Columbia, Duke-Fuqua, ETH Zurich, Florida International, IFPR, Penn, Santa Fe Institute, Toronto, UCLA, UCSD, UC Davis, USC, USC-Marshall, Western Ontario, and Yale; and conference participants at the GTS 3rd World Congress, SWET 2008, the 2008 Stony Brook Workshops, the 2009 NBER Organizational Economics Working Group, and NAWMES 2011. Jacob Johnson, Jong-Myun Moon, and Aniela Pietrasz provided excellent research assistance. Miller thanks Yale and the Cowles Foundation for hospitality and financial support, and Partha Dasgupta for inspiration to pursue this topic. Watson thanks NSF (SES-0095207), NOAA Fisheries Service, and the Cowles Foundation at Yale for hospitality and financial support.

Contents

1 Introduction

Many economic relationships (such as partnership, employment, and buyer-supplier relationships) are ongoing and governed in whole or in part by self-enforced incentives. The relational contracts literature studies these relationships using the framework of infinitely repeated games, but repeated games suffer from a vast multitude of equilibria, particularly when players are patient.[1] So equilibrium theory alone offers little hope for predicting behavior, or for identifying parameters from observed behavior.

To understand these ongoing relationships, we need a theory of equilibrium selection. Informally, it seems that players in an ongoing relationship must coordinate closely in order to select "their" equilibrium from the multitude, so it often makes sense to suppose that they coordinate on a Pareto-optimal equilibrium. This intuition underlies many classic results in the relational contracts literature. But this intuition raises two issues. First, payoffs on the Pareto frontier may be supported by the threat of punishments that depart from the Pareto frontier. What if the players can re-coordinate their continuation play to escape from a Pareto-dominated punishment? Second, there are typically many equilibria on the Pareto frontier, where one player's gain is necessarily a loss for others. How do they decide which Pareto-optimal equilibrium to select? Though these issues have arisen in the relational contracts literature, so far a unified solution has been lacking.[2]

We solve these issues by modeling equilibrium selection as a noncooperative bargaining process embedded in a repeated game with transfers. In each period, the players engage in cheap-talk bargaining (a generalized alternating-offer protocol) and make voluntary transfers prior to playing the stage game. We propose three axioms that endow cheap-talk messages with endogenous meaning. We show that any equilibrium satisfying the three axioms—which we call a *contractual equilibrium*—has a simple canonical representation. In a contractual equilibrium, the bargaining protocol influences the distribution of welfare between the players and also the level of welfare they can jointly attain. These welfare implications have a simple representation in terms of relative bargaining power, which is derived from the details of the bargaining protocol.

We provide an explicit algorithm for constructing contractual equilibria, and we characterize their efficiency and allocative properties. We also extend the theory to games with more than two players, imperfect public monitoring, and heterogeneous discount factors. To illustrate the utility

[1]Viz.: the folk theorem (Fudenberg and Maskin 1986; Fudenberg, Levine, and Maskin 1994, etc.).

[2]There is also the renegotiation-proofness literature, which addresses the first issue but generally not the second. We discuss renegotiation proofness in some depth later in this section.

of contractual equilibrium for the study of relational contracts, we apply it to a canonical principal-agent problem with moral hazard (following Levin 2003); we also use contractual equilibrium to address the question of whether external enforcement crowds out self enforcement (motivated by Baker, Gibbons, and Murphy 1994).

Axioms on endogenous meaning The three axioms constitute the conceptual core of our theory. The first axiom, *Internal Agreement Consistency (IAC)*, specifies that if an offer is proposed and accepted to play as if switching to a different history in the same equilibrium, then the "agreement" has meaning: the players play as agreed. The idea is that the continuation play associated with such a deviation constitutes an equilibrium, and so the players recognize in the bargaining phase that it is a valid option for them to agree on.[3] By itself, IAC does not change the set of equilibrium payoffs in the game, as shown by Theorem 1. A deviant agreement to switch histories can always be discouraged by punishing the player who proposes it and rewarding the player who rejects it.

The second axiom, *No-Fault Disagreement (NFD)*, embodies a theory of disagreement. It specifies that if the players do not reach an agreement in a given period, then their continuation play should not depend on how bargaining broke down. The idea is that no player should be selectively punished for putting an innovative offer on the table, or for rejecting the equilibrium offer in hopes of being able to make an innovative counteroffer. Also, disagreement implies zero monetary transfers in the current period, which is equivalent to saying that failure to make a required transfer induces disagreement. NFD does allow that continuation play under disagreement may be Pareto dominated by continuation play under agreement—indeed, this property plays an important role in our results. But by itself NFD has little influence over the set of equilibrium payoffs in the game (Theorem 2). Specifically, it cannot eliminate any payoffs on the Pareto frontier of what can be attained by pure strategy subgame perfect equilibria with transfers.

Putting IAC and NFD together, however, leads to a simple characterization. By NFD, for each history there is a well-defined payoff vector associated with disagreement. By IAC, every payoff available under agreement at any history is available at every history. Together, these properties endow the bargaining phase at any given history, viewed as a game in itself, with a unique subgame perfect equilibrium outcome. As the bargaining literature has established, the players agree on a payoff vector in the bargaining set that splits the surplus, relative to the disagreement

[3] IAC has a similar flavor to the notions of "internal consistency" of Bernheim and Ray (1989) and "weak renegotiation proofness" of Farrell and Maskin (1989), but is significantly weaker. It merely makes certain deviant agreements available, without assuming that players must accept them.

payoff, according to the allocation of bargaining power that arises from the bargaining protocol. For the repeated game, now applying the dynamic programming techniques of Abreu, Pearce, and Stacchetti (1990) yields a set of disagreement payoffs that are enforceable with respect to the set of agreement payoffs, which in turn are supported as bargaining outcomes with respect to the disagreement payoffs. The set of equilibrium payoffs available under IAC and NFD are typically much smaller than the set of all subgame perfect payoffs. For instance, in some Prisoners' Dilemma games, the only possible outcome is infinite repetition of the stage game equilibrium (see Example 1). Generally, though, the payoffs available under IAC and NFD can be partially Pareto ranked.

The third axiom, *Pareto External Agreement Consistency (PEAC)*, formalizes the intuition that players should select a Pareto optimal equilibrium. PEAC specifies that if an offer is proposed and accepted to play as if switching to an equilibrium that is *fully Pareto dominant* (Definition 1) among those that satisfy IAC and NFD, then the agreement has meaning and the players follow through on it. The idea is that the continuation play associated with such a deviation constitutes a valid equilibrium, so the players recognize in the bargaining phase that it is a valid option for them to agree on.[4] Under PEAC, the players agree to play a Pareto-dominant subgame perfect equilibrium among those that satisfy IAC and NFD.

Characterization of contractual equilibrium Our representation theorem, Theorem 3, shows that the set of contractual equilibria has a simple dynamic programming representation. In this dynamic program, the bargaining and transfer phases are replaced with their backward induction solution, in which the players share the surplus (measured relative to the value of disagreement) in proportions that depend on the details of the alternating-offer bargaining protocol. We characterize the set of contractual equilibrium values, and show that it is small enough to make a number of sharp predictions. Which payoff is selected from the this set depends on what would happen if the players disagreed in every period. Since coordination arises from agreement in bargaining, if the players never agree then they are not actively coordinating. Their initial selection thus depends on whatever is the exogenously given custom or status quo that they fall back on when they continually fail to agree.

Theorem 4 shows that, for any discount factor, there is a non-empty set of contractual equilibrium values, and it is a compact line segment of slope -1. Theorem 5 provides an explicit algorithm for constructing this set. The algorithm involves computing two optimal disagree-

[4]PEAC has a similar flavor to the Bernheim and Ray (1989) notion of "external consistency;" but, like IAC, it merely makes certain deviant agreements available, without assuming that players must accept them.

ment points that pin down the endpoints of the line segment. Section 5.2 provides necessary and sufficient conditions for patient players to attain efficiency, as well as several simple sufficient conditions. Section 5.3 discusses the role of relative bargaining power as determined by the bargaining protocol. Specifically, Theorem 9 shows that the sum of the players' payoffs is maximized by assigning all bargaining power (i.e., a monopoly over the right to make offers) to one player or the other. Even in symmetric games, asymmetric bargaining power can be strictly more efficient than symmetric bargaining power (see Example 2). Section 6 generalizes contractual equilibrium to games with more than two players, imperfect public monitoring, and heterogeneous discount factors. Theorem 10 shows that the set of contractual equilibrium values, when measured in total discounted payoffs, forms a non-empty, compact hyperpolygon normal to the vector of ones. Theorem 11 provides an explicit constructive algorithm for two-player games with imperfect public monitoring and heterogeneous discount factors. Finally, Section 7 applies contractual equilibrium to two seminal models from the relational contracts literature. We discuss these two models and the wider relational contracts literature in more detail after first addressing another related literature—renegotiation proofness.

Contrast with renegotiation proofness The renegotiation-proofness literature addresses the problem of Pareto-dominated continuation play by ruling it out. That is, for a given equilibrium, if the continuation from any particular history is dominated by a qualified alternative, then the equilibrium is removed from consideration. Different notions of renegotiation proofness arise from different restrictions on the class from which Pareto-dominant alternatives may be drawn. The study of renegotiation proofness in infinitely repeated games was initiated by Rubinstein (1979), Bernheim and Ray (1989), Farrell and Maskin (1989), Asheim (1991), and Pearce (1987).[5] Renegotiation proofness does not model renegotiation explicitly. Parts of the literature do contemplate the possibility of renegotiation in the form of bargaining, but merely as a conceptual device to rule out equilibria.

By contrast, contractual equilibrium entails an explicit account of the bargaining in each period, and it posits a theory of what happens under disagreement. Because the set of possible disagreement continuation values is generally a strict subset of the possible agreement values, the players' bargaining powers play a role in determining the agreement in each period. Importantly, under disagreement the continuation value may lie below the relevant Pareto frontier. Such an inferior continuation is not contemplated in the renegotiation-proofness concepts.

[5] Renegotiation proofness is further developed by Abreu, Pearce, and Stacchetti (1993), Abreu and Pearce (1991), Bergin and MacLeod (1993), Blume (1994), Ray (1994), and Asheim (1997), among others.

It is helpful to note variations in the renegotiation-proofness literature on the timing of transfers and bargaining within a period. We consider bargaining followed by voluntary transfers and then the stage game (which itself may involve transfers). For renegotiation proofness in repeated games with transfers, Baliga and Evans (2000) study transfers that occur simultaneously with actions, while Kranz and Ohlendorf (2010) and Fong and Surti (2009) study transfers that occur in a separate phase. Kranz and Ohlendorf look at variations of renegotiation proofness in which it is applied either only before transfers or both before transfers and stage-game actions in a period.

To appreciate the central role that bargaining power plays in our analysis, consider the notion of renegotiation proofness that is closest to contractual equilibrium: the application of "strong optimality" (Levin 2003) only before the transfer phase in a period. Kranz and Ohlendorf show that this notion does not refine the Pareto frontier of subgame perfect equilibrium payoffs. A deviator in the transfer phase is punished by receiving his worst enforceable action-phase payoff, both in the immediate action phase and in the continuation game. In contrast, under contractual equilibrium the continuation payoff for the next period is determined by bargaining, so unless the deviator has no bargaining power, under agreement he will always obtain a payoff strictly higher than his worst enforceable action-phase payoff. Thus contractual equilibrium affects the Pareto frontier of attainable payoffs, in a way that is consistent with the players' relative bargaining power.

On the question of whether a renegotiation-proofness concept should be applied everywhere within a period, our view is that this sidesteps the key objective of providing an explicit account of bargaining. Once bargaining is described in the extensive form, it is clear that disagreement must always be a feasible outcome within a period. The implications of bargaining then spring from assumptions made about disagreement. Our modeling exercise seeks to highlight this point.

Relational contracts The relational contracts literature examines how agreements (such as in an employment relationship) can be self-enforced through repeated play. Some approaches in this literature do not analyze bargaining or negotiation at all, or allow bargaining only at the outset of the relationship.[6] Other approaches allow for renegotiation once the relationship is underway, but assume that the parties permanently separate if they disagree, or at least switch permanently to a stage game equilibrium or a spot contract (e.g., Hermalin, Li, and Naughton 2011; Jackson and Palfrey 1998; Levin 2002; MacLeod and Malcomson 1989, 1998; Schmidt and Schnitzer 1995). The assumption that disagreement leads to separation is also commonly used to integrate wage

[6]For example, Baker, Gibbons, and Murphy (1994, 2002); Chassang (2010); Doornik (2006); Fuchs (2007); Kvaløy and Olsen (2006); Levin (2002); Pearce and Stacchetti (1998); Radner (1985); Ramey and Watson (2001, 1997); Rayo (2007); Schmidt and Schnitzer (1995); Thomas and Worrall (1988).

bargaining into macroeconomic models with labor market search frictions.[7]

Most closely related to our work are those papers that allow for continued interaction following disagreement, and in which bargaining power plays a role over time. However, the prior literature makes strong assumptions on what the players may do under disagreement. These papers impose a theory of disagreement in which, until their next agreement, the players either receive an exogenously specified, temporary outside option (Fong and Li 2010a,b; Halac 2010; Ramey and Watson 2002); or play a stage-game Nash equilibrium (Klimenko, Ramey, and Watson 2008; Levin 2003; Raff and Schmidt 2000).

In Section 7.1, we examine contractual equilibria of the principal-agent model studied by Levin (2003). In Levin's model, the principal makes a take-it-or-leave-it offer to the agent in every period. For most of his analysis, Levin assumes that the parties play the stage game equilibrium if the agent rejects the principal's offer. But Levin also considers a variation on renegotiation proofness called "strong optimality," in which continuation play from the start of each period is always on the Pareto frontier, and shows that any optimal contract can also be made strongly optimal. This result arises from the observation that any unexpected offer from the principal can be interpreted by the agent as a signal to revert to the stage game equilibrium for the period. This is why bargaining power plays no role in Levin's analysis—the equilibria he studies violate IAC. The principal cannot make a meaningful deviant offer, so she cannot benefit from her monopoly over making proposals.

Our analysis of the model, in contrast, shows that in contractual equilibrium the agent's effort is increasing in his own bargaining power, and is zero if he has no bargaining power at all. Indeed, the principal's ideal level of bargaining power is intermediate—if she has all the bargaining power, she cannot commit to payments that would motivate the agent; if she has no bargaining power then the agent extracts all the surplus.

In Section 7.2 we apply contractual equilibrium to the interaction of "explicit" (externally enforced, such as by a court) contracts and "implicit" (self-enforced) contracts, following Baker, Gibbons, and Murphy (1994). Both the analysis of Baker, Gibbons, and Murphy and related work by Schmidt and Schnitzer (1995) focus on short-term externally enforced contracts, to which the parties revert if they fail to agree on a self-enforced contract. The availability of a short-term externally enforced contract that is better than no contract at all constrains the parties' ability to

[7] For example, Cahuc, Postel-Vinay, and Robin (2006); den Haan, Ramey, and Watson (2000); Diamond (1982); Diamond and Maskin (1979); Flinn (2006); Hosios (1990); Mortensen and Pissarides (1994); Moscarini (2005); Pissarides (1985, 1987). Den Haan, Ramey, and Watson (2003) and Genicot and Ray (2006) apply a similar approach to credit markets.

punish deviations, and therefore can lead to lower payoffs in equilibrium.

Contractual equilibrium, in contrast, allows the parties to reevaluate their entire relationship whenever they engage in bargaining, and does not bind them to consider only short-term externally enforced contracts. Still, we can use contractual equilibrium to study the interaction between self-enforced contracts and *long-term* externally enforced contracts. A long-term externally enforced contract is signed at the outset of the relationship and stays in legal force unless the parties mutually agree to suspend or terminate it. In particular, it stays in force when they disagree. We show that an optimal long-term externally enforced contract induces the parties to waste resources when they disagree, in order to raise the stakes of bargaining. With higher stakes, the difference between the agent's reward and punishment states rises, yielding higher-powered incentives and therefore higher effort. In short, being able to design an explicit contract unambiguously makes the parties better off.

2 Repeated games with bargaining

Consider a standard two-player repeated game, augmented with cheap-talk bargaining and transferable utility. In each period, before the players play the stage game, they bargain according to a generalized alternating-offer protocol. The players can make voluntary transfers before the stage game. The stage game can also be defined to include voluntary transfers that occur simultaneously with other actions. There is no external enforcement of any kind.

2.1 Extensive form

Formally, a two-player game in this class is defined by a stage game $\langle \mathcal{A}, u \rangle$, a discount factor $\delta \in (0, 1)$ by which the players exponentially discount their payoffs across periods, and a bargaining protocol that we will describe shortly. Here $\mathcal{A} \equiv \mathcal{A}_1 \times \mathcal{A}_2$ is the space of action profiles and $u \colon \mathcal{A} \to \mathbb{R}^2$ is the stage-game payoff function. We express repeated game utilities in discounted average terms to facilitate comparison to stage game payoffs.

Each period comprises four phases: (1) the public randomization phase, (2) the bargaining phase, (3) the transfer phase, and (4) the action phase. In the public randomization phase, the players observe an arbitrary public randomization device.[8] In the transfer phase, the players simultaneously make voluntary, non-negative monetary transfers; that is, each player decides how

[8]Such devices are standard in repeated games analysis (see Fudenberg and Maskin 1986). Here, in particular, the randomization device should be viewed as having been constructed by the players in their process of equilibrium selection.

much money to pay to the other, where money enters their utility quasilinearly. In the action phase, the players play the stage game $\langle \mathcal{A}, u \rangle$.

In the bargaining phase, the players engage in a generalized alternating-offer bargaining protocol. The number of potential rounds of bargaining may be finite or infinite, but all these rounds occur in a mere instant, so bargaining does not delay the later phases of the period. (If bargaining involved delay, it would not be cheap talk.) There is an exogenous random recognition process $\rho \in (\Delta\{1,2\})^\infty$ that selects one of the players to make a verbal statement in each round of bargaining, where $\rho_{i,\ell}$ is the probability that player i is recognized in round ℓ. The selected player—called the *proposer*—selects a *proposal* from some language set $\mathcal{L}$ and this is observed by the other player, who is called the *responder*. We assume $\mathcal{L}$ contains a special "no offer" element.

If the proposal is something other than "no offer" then the responder chooses a *response* from $\{\text{yes}, \text{no}\}$; otherwise, the responder has no action. If the response is "yes," indicating the responder's acceptance, then the bargaining phase ends and the game proceeds to the transfer phase. If the response is "no" or if the proposal was "no offer," then bargaining may either break down or continue to another round. If bargaining breaks down then play proceeds to the transfer phase. Breakdown is triggered randomly by a process $\beta \in [0,1]^\infty$ with $\prod_{\ell=1}^{\infty}(1-\beta_\ell) = 0$, where β_ℓ is the probability that breakdown occurs after a "no offer" proposal or a "no" response in bargaining round ℓ. Both the recognition process and the breakdown process are invariant to the time period and the history; they depend only on the round of bargaining within a period.

The language $\mathcal{L}$ should be sufficiently large that each player can use it to propose to the other how to coordinate their continuation play.[9] Specifically, we assume that $\mathcal{L}$ contains the space of possible continuation payoff vectors from the voluntary transfer phase.

Assumption 1 (Rich language)**.** $\mathbb{R}^2 \subset \mathcal{L}$.

We think of a proposal $w \in \mathbb{R}^2$ as a suggestion to coordinate on play in the continuation game to achieve w as the average payoffs in the continuation game.

Next we develop some notation for histories. We start by describing outcomes of the bargaining phase within a period. Suppose bargaining lasts ℓ rounds, at which point an offer is accepted or there is exogenous breakdown. Such an outcome is fully described by an element of $(\{1,2\} \times \mathcal{L})^\ell \times \{\text{agreement}, \text{breakdown}\}$, where, for each round of bargaining, $\{1,2\}$ accounts

[9]One way to ensure that the language is large enough would be to assume that $\mathcal{L}$ contains descriptions of all continuation strategy profiles in the game. However, since such a construction would be circular (strategies would specify history-dependent statements in each negotiation phase, and these statements would include the description of an entire strategy profile), it would lead to a complicated, technical issue regarding whether an appropriate "universal language" exists.

for the identity of the proposer and $\mathcal{L}$ for the proposal. For the first $\ell - 1$ rounds, either the proposal was "no offer" or the response was "no." Then, for round ℓ, "agreement" signifies that the last proposal was accepted, whereas "breakdown" signifies that bargaining broke down.[10] Thus the set of possible bargaining outcomes is

$$\mathcal{B} \equiv \bigcup_{\ell=1}^{\infty} (\{1, 2\} \times \mathcal{L})^{\ell} \times \{\text{agreement, breakdown}\}. \tag{1}$$

The set of full-period histories, including the "null history," is

$$\mathcal{H} \equiv \bigcup_{t=0}^{\infty} \big(\Omega \times \mathcal{B} \times \mathbb{R}_+ \times \mathbb{R}_+ \times \mathcal{A}\big)^t, \tag{2}$$

where Ω is the state space of the arbitrary public randomization device. Here the first $\mathbb{R}_+$ accounts for transfers from player 1 to player 2, and the second $\mathbb{R}_+$ accounts for transfers from player 2 to player 1. The net of these transfers is summarized by a transfer vector in $\mathbb{R}^2_0 \equiv \{m \in \mathbb{R}^2 : m_2 = -m_1\}$.

A *history to the transfer phase* is an element of $\mathcal{H} \times \Omega \times \mathcal{B}$. A *history to the action phase* is an element of $\mathcal{H} \times \Omega \times \mathcal{B} \times \mathbb{R}_+ \times \mathbb{R}_+$. A history to the transfer phase or to the action phase is *under agreement* if the just-completed bargaining phase ended with "agreement"; otherwise the history is *under disagreement*.

2.2 Recursive characterization of equilibrium values

We characterize the set of subgame perfect equilibrium payoffs by extending the recursive methods of Abreu, Pearce, and Stacchetti (1990), working backward through the phases of a given period. Consider first the action phase. Let $\Delta^{\mathrm{U}}\mathcal{A}$ be the set of probability distributions over $\mathcal{A}$ that are uncorrelated across dimensions; i.e., the set of mixed action profiles that can arise from independent mixed actions for each player. In equilibrium, the mixed action profile must constitute a Nash equilibrium, taking the players' continuation play from the following period as given. Formally, if $g\colon \mathcal{A} \to \mathbb{R}^2$ gives the continuation value as a function of the realized action profile $a \in \mathcal{A}$, and if $\alpha \in \Delta^{\mathrm{U}}\mathcal{A}$ is a Nash equilibrium of $\langle \mathcal{A}, (1-\delta)u + \delta g \rangle$, then we say *$g$ enforces α*. In

[10] Technically, it is not feasible for the outcome to entail "agreement" if the last proposal is "no offer." One can either remove such sequences from consideration or just assume that a sequence ending in (no offer, agreement) is equivalent to one ending in (no offer, breakdown).

this case, the continuation value from the action phase is

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

If $W \subset \mathbb{R}^2$ is the set of feasible continuation values from the beginning of the next period, then the set of continuation values that can be supported from the action phase of the current period is:

$$D(W) \equiv \left\{ w \in \mathbb{R}^2 \,\middle|\, \exists\, g\colon \mathcal{A} \to W \text{ and } \alpha \in \Delta^{\mathrm{U}}\mathcal{A} \text{ s.t. } g \text{ enforces } \alpha \text{ and Eq. 3 holds} \right\}. \tag{4}$$

Define $\overline{D}_i(W) \equiv \cup_{w \in D(W)} [w_i, \infty)$ and let $\overline{D}(W) \equiv \overline{D}_1(W) \times \overline{D}_2(W)$. As shown by Goldluecke and Kranz (2010), the set of continuation values supported from the transfer phase is:

$$C(W) \equiv \left\{ w \in \mathbb{R}^2 \,\middle|\, \exists\, m \in \mathbb{R}^2_0 \text{ and } w' \in D(W) \text{ s.t. } w = (1-\delta)m + w' \text{ and } w \in \overline{D}(W) \right\}. \tag{5}$$

To understand how values in $C(W)$ can be achieved, consider any w, m, and w' that satisfy the conditions of Eq. 5. Without loss of generality, take the case of $m_1 \geq 0$ so that player 2 is supposed to make a transfer to player 1. Prescribe that player 2 make this transfer and then the players continue with behavior to achieve w' from the action phase. If player 2 does not make the required transfer then the players coordinate on behavior to achieve a value $\underline{w} \in D(W)$ from the action phase, where $\underline{w}$ is selected so that $w_2 \geq \underline{w}_2$. Clearly the prescribed behavior is rational from the transfer phase.[11]

Operator C provides the basis for the recursive formulation of subgame-perfect equilibrium payoffs. Given a set of equilibria S, let $V(S)$ be the set of continuation values attained by S starting from any full-period history. When V is evaluated at a singleton set $\{s\}$, we call it the *value set* of s, and abuse notation to write $V(s)$. Let S_{SPE} be the set of all subgame-perfect equilibria of our model with bargaining. From the construction above, it is clear that $V(S_{\mathrm{SPE}})$ is the largest fixed point of co $C(\cdot)$, where "co" denotes the convex hull. The convex hull here accounts for public randomization at the beginning of each period. Also, observe that $V(S_{\mathrm{SPE}})$ is identical to the set of subgame perfect equilibrium payoffs in an otherwise-identical game with no bargaining phase.[12]

[11]Nothing more can be supported since feasibility requires $w = (1-\delta)m + w'$ for some $m \in \mathbb{R}^2_0$ and $w' \in D(W)$, and each player i can guarantee himself some value in $\overline{D}(W)$ by making a zero transfer. Also note that it is not helpful to have both players make positive transfers. From such a specification, one can reduce the transfers equally until one of the transfers is zero. This does not affect the continuation value but loosens the incentive constraints.

[12]The players can always ignore the bargaining phase, so $V(S_{\mathrm{SPE}})$ includes all no-bargaining payoffs. Similarly, since

3 Axiomatization of endogenous meaning

So far, statements in $\mathcal{L}$ have neither any particular meaning nor any effect on the set of equilibrium payoffs. We next develop axioms that refine the equilibrium set by imposing meaning in equilibrium. Specifically, a proposal $w \in \mathbb{R}^2$ expresses the suggestion that the players should coordinate their behavior to achieve the continuation value w from the transfer phase in the current period.

3.1 A foundation for meaningful proposals

In an equilibrium $s \in S_{\text{SPE}}$, the players recognize that any payoff in $C(\text{co } V(s))$ is attainable from the transfer phase, by randomizing over continuation values achieved by s using the public randomization device at the beginning of the next period. We posit that the players should have the opportunity to agree upon such a payoff and then follow through on their agreement. That is, any proposal to play as if conditionally switching to different histories in the same equilibrium should be taken seriously. Our first axiom requires that, on or off the equilibrium path, if there is agreement on a continuation value w that is consistent with the equilibrium strategy profile (the proposer offers $w \in C(\text{co } V(s))$ and the responder says "yes"), then the players proceed in a way that achieves value w.

Axiom IAC (Internal agreement consistency)**.** For every history to the transfer phase under agreement, if the agreement is of the form $w \in C(\text{co } V(s))$ then continuation play yields the value w.

Let S_{IAC} be the subset of S_{SPE} that satisfies IAC. Note that this axiom does not require that the players reach agreement in equilibrium. It merely defines the meaning of a class of agreements, which refines the equilibrium set due to the consequences in off-equilibrium-path contingencies. But although IAC eliminates some equilibria, it does not on its own refine the set of equilibrium values.

Theorem 1 (IAC has no bite)**.** $V(S_{\text{IAC}}) = V(S_{\text{SPE}})$.

The proof, in Appendix A.1, actually strengthens the result by invoking a stronger consistency axiom that also addresses agreements to switch to continuations that are supported by other equilibria. Here is the basic intuition for the proof: Under IAC there may be deviant proposals that a

both games feature arbitrary public correlation devices and the bargaining phase is pure cheap talk, every $v \in V(S_{\text{SPE}})$ is attainable without bargaining.

proposer can offer, which, if implemented, would make him better off than if he made his equilibrium proposal. However, IAC does not constrain continuation play after such a deviant proposal is rejected by the responder. Given any equilibrium $s \in S_{\text{SPE}}$, we construct an alternative equilibrium $s' \in S_{\text{IAC}}$ as follows. From any given history to the bargaining phase, the players are supposed to propose and accept the continuation value that strategy profile s would support. If a deviant proposal is accepted, then play in the continuation is designed to conform to IAC. If the responder says "no," however, then the players coordinate to achieve a continuation value that punishes the proposer and rewards the responder. In this way, the proposer is deterred from making a deviant proposal.

As the proof of the theorem makes clear, the agreement consistency axiom is undermined by the fact that continuation play can vary in arbitrary ways following a disagreement. In order to narrow the set of equilibrium values, we need a theory of disagreement.

3.2 A no-fault theory of disagreement

We next add a *no-fault disagreement* axiom embodying a theory of how players behave under disagreement. It requires that continuation play in the absence of agreement—though it may vary with the history of actions, transfers, and past agreements—should not vary with the manner in which bargaining broke down.

Disagreement can arise in myriad ways. Bargaining may randomly break down after a proposer says "no offer" or a responder says "no", following any sequence of equilibrium or deviant offers. That disagreement play should not be sensitive to the manner of breakdown represents the nostrum that "nothing is agreed until everything is agreed." That is, no individual player can be held personally responsible for a breakdown of bargaining. In particular, no player can be singled out for punishment for making a deviant offer that is rejected, or for rejecting an offer in hopes of being able to make a counteroffer. We also assume that play from the action phase under disagreement is insensitive to the transfers just made. This captures the idea that a player essentially voids an agreement by refusing to make the required up-front payment.[13]

[13]An alternative interpretation is that a deviation in the transfer phase triggers a disagreement as well; i.e., transfers "seal the deal." This interpretation gives any player called upon to make a transfer one last chance to cancel the agreement, essentially by refusing to sign the check. In the realm of relational contracts, where there is always at least a rudimentary legal system in the background, this interpretation embodies the fact that signing a contract can create a legal obligation to make an unconditioned payment, even when no other contract clauses are legally enforceable. Adopting this interpretation formally would require extra notation to keep track of the transfer required in agreement, and we would have a slightly different version of IAC. We obtain the same implications with the assumption that continuation play is insensitive to transfers under disagreement, as a player always has the option of forcing disagreement and making no transfer.

Axiom NFD (No-fault disagreement)**.** For every history to the action phase under disagreement, continuation play is measurable with respect to the full-period history at the end of the previous period and the realization of the public randomization device in the current period.

Let S_{NFD} be the subset of S_{SPE} that satisfies NFD. A variant of operator C can be used to characterize set $V(S_{\text{NFD}})$:

$$\hat{C}(W) \equiv \left\{ w \in \mathbb{R}^2 \,\middle|\, \exists\, m \in \mathbb{R}^2_0 \text{ and } \underline{w}, w' \in D(W) \text{ s.t. } w = (1-\delta)m + w' \text{ and } w \geq \underline{w} \right\}. \tag{6}$$

Note that $\hat{C}(W) \subseteq C(W)$ because $D(W) \subset \overline{D}(W)$. Under agreement, the players make a transfer m and then attain a continuation value of w' from the action phase. The point $\underline{w} \in D(V(s))$ gives the continuation value from the action phase under disagreement; each player i can guarantee himself a payoff of at least $\underline{w}_i$ by forcing disagreement in the bargaining phase (always proposing "no offer" and saying "no") and making no transfer. Since the transfer makes both players better off than under disagreement, it is incentive compatible. Since $w' \in D(W)$, the actions are incentive compatible as well. Thus $V(S_{\text{NFD}})$ is the largest fixed point of co $\hat{C}(\cdot)$.

Although NFD constrains the meaning of disagreements, by itself NFD does little to constrain the set of attainable payoffs. In particular, in the realm of pure-strategy equilibria, NFD does not alter the Pareto boundary of equilibrium continuation values. Let $S^{\text{ps}}_{\text{SPE}}$ be the set of pure-strategy subgame-perfect equilibria of the model, and let $S^{\text{ps}}_{\text{NFD}}$ be the subset of equilibria that satisfy NFD. For any set of payoffs $W \subset \mathbb{R}^2$, let $P(W)$ be the Pareto frontier of W.

Theorem 2 (NFD has little bite)**.** $P(V(S^{\text{ps}}_{\text{SPE}})) = P(V(S^{\text{ps}}_{\text{NFD}})) \subseteq V(S_{\text{NFD}})$.

The proof, in Appendix A.2, shows a stronger result: any payoff attainable without transfers—regardless of whether it is on the Pareto frontier—is also contained in $V(S_{\text{NFD}})$.

In combination with IAC, however, NFD has strong effects. In particular, it rules out the kinds of punishments used to discourage deviant proposals in the proof of Theorem 1. Following any history, the bargaining phase (viewed in isolation as a game in itself) has a unique disagreement outcome by NFD, and a non-degenerate bargaining set by IAC. Therefore it has unique subgame-perfect equilibrium outcome that depends only on the recognition process, the breakdown process, the continuation value under disagreement, and the joint surplus attainable under agreement. Moreover, in equilibrium the players share the surplus in proportions that do not depend on the magnitude of the surplus (see Binmore 1987; Rubinstein 1982). We characterize the implications of NFD and IAC formally in Section 4.

Although the combination of NFD and IAC eliminates many equilibria, generally multiple Pareto-ranked equilibrium payoffs still survive. Our third axiom allows the players to meaningfully discuss whether to switch to an equilibrium that is "better" than the one they are currently playing. The following definition formalizes what makes one equilibrium better than another.

Definition 1. A payoff set $W' \subset \mathbb{R}^2$ *fully Pareto dominates* another payoff set W if (i) for every $v' \in W'$ there is no $v \in W \setminus W'$ satisfying $v \geq v'$ (component-wise), and (ii) for every $v \in W$ there exists $v' \in W'$ such that $v' \geq v$ (component-wise).

We say that a strategy profile s' *fully Pareto dominates* another strategy profile s if co $V(s')$ fully Pareto dominates co $V(s)$. Because full Pareto dominance is a demanding notion of dominance, the following axiom—which requires the players to honor any agreement to deviate to an equilibrium that fully Pareto dominates their current equilibrium—is relatively weak.

Axiom PEAC (Pareto external agreement consistency)**.** For every history to the action phase under agreement in equilibrium s, if the agreement takes the form $w \in C(\text{co } V(s'))$, where $s' \in S_{\text{IAC}} \cap S_{\text{NFD}}$ and co $V(s')$ fully Pareto dominates co $V(s)$, then continuation play yields the value w.

We define a subgame perfect equilibrium to be a *contractual equilibrium* if it satisfies IAC, NFD, and PEAC. Accordingly, let S_{CE} be the subset of $S_{\text{IAC}} \cap S_{\text{NFD}}$ that satisfies PEAC. The set of contractual equilibrium values has a sharp representation: $V(S_{\text{CE}})$ is identical to a "dominant bargaining self-generated set" that we identify in the next section. The rest of this paper is devoted to characterizing these sets.

4 Representation

To characterize contractual equilibrium values, we compare the bargaining phase in an individual period to a simple bargaining game in which (i) the players bargain directly over a fixed set of payoff vectors and (ii) the generalized alternating-offer protocol is of the form described in Section 2.1. Analysis of this simple bargaining game yields an operator that we embed in a recursive construction of continuation values.

In the simple bargaining game, the players bargain over the selection of a payoff vector from a set $W \subset \mathbb{R}^2$. If they fail to reach an agreement, then the payoff vector is some value $\underline{w} \in W$. Suppose that the sum of the players' continuation payoffs — which we call the *welfare level* — is constant along the Pareto boundary of W, and suppose this boundary includes all points above $\underline{w}$

with the same welfare. Then, in expectation, the subgame perfect equilibrium payoff vector w is unique and satisfies:

$$w_1 + w_2 = \max_{w' \in W} \big(w'_1 + w'_2\big) \quad \text{and} \quad w = \underline{w} + \pi\big(w_1 + w_2 - \underline{w}_1 - \underline{w}_2\big), \tag{7}$$

where for each player i, $\pi_i = \sum_{\ell=1}^{\infty}\big(\rho_{i,\ell}\beta_\ell \prod_{k=1}^{\ell-1}(1-\beta_k)\big)$. In other words, the players maximize welfare and they split the surplus in fractions π_1 and π_2. This conclusion is stated in Lemma 6 in Appendix A.3. We shall write $\pi = (\pi_1, \pi_2)$; note that $\pi_1 + \pi_2 = 1$.

Next we construct a set of payoff vectors by calculating the equilibrium payoffs that can arise in the simple bargaining game when $\underline{w}$ can be any disagreement value in $\underline{W}$:

$$B(W, \underline{W}) \equiv \left\{ \underline{w} + \pi\Big(\max_{w' \in W}(w_1 + w_2) - \underline{w}_1 - \underline{w}_2\Big) \;\Big|\; \underline{w} \in \underline{W} \right\}. \tag{8}$$

Returning to the repeated game, recall that operators D and C produce the sets of continuation values from the action phase and transfer phase of a given period, respectively, as a function of a set of continuation values W from the start of the following period. Let us interpret $C(W)$ as the set of payoff-vector alternatives for the players in the bargaining phase. Further, interpret $D(W)$ as the set of disagreement payoffs, which are attainable from the transfer phase with transfers of zero. The bargaining operator, along with public randomization at the beginning of the period, yields the compound operator $\operatorname{co} B(C(\cdot), D(\cdot))$.[14] We examine fixed points of this compound operator.

Definition 2. A non-empty set $W \subset \mathbb{R}^2$ is called a *bargaining self-generated (BSG)* set if $W = \operatorname{co} B(C(W), D(W))$.

Our next definition combines bargaining self-generation with full Pareto dominance.

Definition 3. A set $W^* \subset \mathbb{R}^2$ is a *dominant BSG* set if it is a BSG set that fully Pareto dominates every other BSG set.

By definition, there can be at most one dominant BSG set. Our "representation theorem" establishes a relation between sets of contractual equilibrium values and bargaining self-generated sets. The theorem shows that the value set of every contractual equilibrium is contained in the dominant BSG set, and for each value in the dominant BSG set there is a contractual equilibrium that delivers it.

[14]Recall that $D(W) \subset C(W)$, welfare is constant along the Pareto boundary of $C(W)$, and $C(W)$ extends to the edges of $D(W)$ along each player's axis. This implies that $B(C(\cdot), D(\cdot))$ is well defined.

Theorem 3 (Representation). *Given a game* $\langle \mathcal{A}, u, \delta, \rho, \beta \rangle$*, let* $\pi_i = \sum_{\ell=1}^{\infty} \left(\rho_{i,\ell} \beta_\ell \prod_{k=1}^{\ell-1} (1 - \beta_k) \right)$ *for each i. Then:*

1. *If W is a BSG set, then* $W \subseteq V(S_{\text{IAC}} \cap S_{\text{NFD}})$*; furthermore, if* $s \in S_{\text{IAC}} \cap S_{\text{NFD}}$*, then* $\text{co}\, V(s)$ *is a BSG set.*

2. $V(S_{\text{CE}})$ *is the dominant BSG set.*

The proof is in Appendix A.3. In principle, every point in $V(S_{\text{CE}})$ is an equally good candidate to be selected as the equilibrium payoff in the first period. By construction, the first period payoff must be $v = B(C(V(S_{\text{CE}})), \{\underline{w}\})$, where $\underline{w} \in D(V(S_{\text{CE}}))$ is the payoff that arises if the players disagree in the first period. But nothing pins down which point in $D(V(S_{\text{CE}}))$ is selected to be the disagreement point in the first period—neither the axioms nor any prior agreement. The same is true of the disagreement payoff in any period following a history in which the players have never agreed. So the problem of selecting a particular equilibrium boils down to determining how the players will play if they have never agreed. In applications, this behavior might arise from custom, the status quo, or social institutions.

Note that the bargaining-protocol parameters ρ and β enter the definition of the compound operator in ways summarized by π_1 and π_2. Thus, for the recursive construction of contractual equilibrium values we can use a *simplified description* of the game, given by $\langle \mathcal{A}, u, \delta, \pi \rangle$. Also, the standard connection between the subgame perfect equilibrium of the non-cooperative bargaining game and a cooperative bargaining solution applies here. In particular, the equilibrium payoff w shown in Eq. 7 coincides with the generalized Nash bargaining solution with weights π_1 and π_2. Thus, we sometimes call π bargaining weights, understanding that they represent the bargaining protocol.

5 Characterization

This section characterizes the set of contractual equilibrium values $V(S_{\text{CE}})$, and provides an algorithm for computing it. The section also contains results on efficiency and how bargaining power affects what values are attainable. The characterization shows that every BSG set (and hence $V(S_{\text{CE}})$) is a line segment with slope -1. Thus, every BSG set has two endpoints $z^1, z^2 \in \mathbb{R}$, satisfying $z_1^1 + z_2^1 = z_1^2 + z_2^2$. Our convention is to define z^1 as the endpoint that favors player 2 ("punishing" player 1) and z^2 as the endpoint that favors player 1. For any such set W, we define

$$\text{span}(W) \equiv z_1^2 - z_1^1 = z_2^1 - z_2^2 \quad \text{and} \quad \text{level}(W) \equiv z_1^1 + z_2^1 = z_1^2 + z_2^2. \tag{9}$$

These values are the vertical distance, or *payoff span*, across W and the joint value, or *welfare level*, of each point in W.

5.1 Existence and construction

We begin by showing that $V(S_{\text{CE}})$ exists for any discount factor.

Theorem 4 (Existence)**.** *For a game $\langle \mathcal{A}, u, \delta, \pi \rangle$, if $\mathcal{A}$ is finite then $V(S_{\text{CE}})$ is a non-empty, compact line segment of slope -1.*

The next result provides a complete, constructive characterization of $V(S_{\text{CE}})$.

Theorem 5 (Construction)**.** *For a game $\langle \mathcal{A}, u, \delta, \pi \rangle$, if $\mathcal{A}$ is finite and $W^* = V(S_{\text{CE}})$, then:*

1. $\text{span}(W^*)$ *is equal to the maximal fixed point of* $\Gamma \equiv \gamma^2 + \gamma^1$*, where for players $i \neq j$,*

$$\begin{aligned} \gamma^j(d) \equiv \max_{\eta,\alpha} & \Big(\pi_j u_i(\alpha) - \pi_i u_j(\alpha) + \frac{\delta}{1-\delta}\eta(\alpha) \Big) \\ \text{s.t.} & \begin{cases} \eta : \mathcal{A} \to [-d, 0], \text{ extended to } \Delta^{\text{U}}\mathcal{A} \\ \alpha \in \Delta^{\text{U}}\mathcal{A} \text{ is a Nash equilibrium of } \big\langle \mathcal{A}_i \times \mathcal{A}_j, (1-\delta)(u_i, u_j) + \delta(\eta, -\eta) \big\rangle; \end{cases} \end{aligned} \tag{10}$$

2. $\text{level}(W^*)$ *is equal to* $\max_{\eta,\alpha} \big(u_1(\alpha) + u_2(\alpha) \big)$

$$\text{s.t.} \begin{cases} \eta : \mathcal{A} \to [0, \text{span}(W^*)], \text{ extended to } \Delta^{\text{U}}\mathcal{A} \\ \alpha \in \Delta^{\text{U}}\mathcal{A} \text{ is a Nash equilibrium of } \big\langle \mathcal{A}, (1-\delta)u + \delta(\eta, -\eta) \big\rangle; \end{cases} \tag{11}$$

3. *The endpoints of W^* are*

$$z^1 = (-1, 1)\, \gamma^1\big(\text{span}(W^*)\big) + \pi\, \text{level}(W^*), \tag{12}$$

$$z^2 = (1, -1)\, \gamma^2\big(\text{span}(W^*)\big) + \pi\, \text{level}(W^*). \tag{13}$$

The rest of this subsection contains the proof of Theorem 4 and Theorem 5. We proceed with a series of lemmas that characterize the geometry of BSG sets and the relation between them; we then identify a dominant BSG set. Proofs of the first three lemmas are presented in Appendix A.4.

Lemma 1. *If $W \subset \mathbb{R}^2$ is a BSG set then it is a bounded, convex subset of a line with slope -1.*

Thus all points in W have the same welfare level. It is possible for a BSG set to not contain one or both of its endpoints. We continue the proof by first constraining attention to BSG sets that are closed—in other words, line segments that contain their endpoints. We argue at the end that any open BSG set is fully Pareto dominated by $V(S_{\text{CE}})$, which is closed.

The next step in the analysis is to characterize the endpoints z^1 and z^2 of an arbitrary closed BSG set. As z^1 and z^2 are extreme points of $B(W, \underline{W})$, there exist disagreement points $\underline{w}^1, \underline{w}^2 \in \underline{W}$ relative to which z^1 and z^2 are the bargaining outcomes satisfying Eq. 7. We will express z^1 and z^2 in relation to optimization problems parameterized by the welfare level and payoff span of W. (Recall that each endpoint is named after the player being punished.)

Lemma 2. *For a given closed BSG set W, $z_1^2 = \pi_1 \operatorname{level}(W) + \gamma^2\big(\operatorname{span}(W)\big)$ and $z_2^1 = \pi_2 \operatorname{level}(W) + \gamma^1\big(\operatorname{span}(W)\big)$. Furthermore, $\gamma^1(d)$ and $\gamma^2(d)$ exist and satisfy $\gamma^2(d) + \gamma^1(d) \geq 0$ for all $d \geq 0$.*

The proof operates by solving the problem of maximizing a player's payoff that can be supported by continuation values in W as a function of the welfare level and payoff span of W. Now we can compare BSG sets by using the functions γ^1 and γ^2. We find that the BSG sets are ranked by full Pareto dominance.

Lemma 3. *Suppose that W is a closed BSG set and W' is any BSG set. If $\operatorname{span}(W) \geq \operatorname{span}(W')$ then W fully Pareto dominates W'.*

The proof follows from the fact that a larger set of continuation values, regardless of its position in $\mathbb{R}^2$, supports a larger range of stage game action profiles.

We next use the functions γ^1 and γ^2 to prove the existence of a fully Pareto dominant BSG set. First, for any BSG set W with endpoints z^1 and z^2,

$$\operatorname{span}(W) = z_1^2 - z_1^1 = z_1^2 - \big(\operatorname{level}(W) - z_2^1\big) = z_1^2 + z_2^1 - \operatorname{level}(W). \tag{14}$$

Substituting for z_1^2 and z_2^1 using Lemma 2 yields $\operatorname{span}(W) = \gamma^2(\operatorname{span}(W)) + \gamma^1(\operatorname{span}(W))$. Thus, the payoff span of a BSG set must be a fixed point of Γ. Further, as the analysis underlying Lemma 2 makes clear, every fixed point of Γ is the payoff span of a BSG set.

Lemma 4. Γ *has a maximal fixed point $d^* \geq 0$.*

Proof. Observe that Γ is non-decreasing because an increase in payoff span relaxes the constraints in the problems that define γ^1 and γ^2. Γ is bounded because u is bounded and δ is fixed. By Tarski's fixed-point theorem, Γ has a maximal fixed point, which is nonnegative because $\Gamma(0) \geq 0$. □

To construct W^* from d^*, let λ^* be the value of $\max_{\eta,\alpha}\big(u_1(\alpha)+u_2(\alpha)\big)$ subject to the constraints in Eq. 11, using d^* in place of span(W). Then construct the endpoints z^1 and z^2 by Eq. 13 using d^* in place of span(W) and λ^* in place of level(W). By construction, $W^* \equiv \text{co}\{z^1, z^2\}$ is a BSG set and its span is maximal over all BSG sets. By Lemma 3, W^* fully Pareto dominates all other closed BSG sets.

To complete the proof, we return to the case of BSG sets that may not be closed. Taking the closure of a BSG set does not necessarily form a BSG set, because new Nash equilibria could emerge in the dynamic program. But for any open BSG set, the following lemma shows that there exists a larger, closed BSG set.

Lemma 5. *If W is a BSG set that is not closed, then there exists a closed BSG set that fully Pareto dominates W.*

Proof. Take the case of a BSG set W that is open on both endpoints. Define $\hat{\gamma}^j(d)$ to be the same as $\gamma^j(d)$ except that (i) its objective is a supremum rather than a maximum, and (ii) the range of η is defined as the open interval $(-d, 0)$ rather than $[-d, 0]$. Then span(W) is a fixed point of $\hat{\gamma}^2 + \hat{\gamma}^1$. By construction, we have $\hat{\gamma}^2 + \hat{\gamma}^1 \leq \gamma^2 + \gamma^1$, so $\Gamma(\text{span}(W)) \geq \text{span}(W)$. Because Γ is increasing and bounded, we thus conclude that $d^* \geq \text{span}(W)$. One can easily confirm that Lemma 2 extends to non-closed BSG sets, which implies that W^* fully Pareto dominates W. Therefore W^* is not open. A parallel argument implies that W^* cannot contain just one of its endpoints. □

As a consequence, W^* fully Pareto dominates all other BSG sets, so we conclude that $W^* = V(S_{\text{CE}})$, proving Theorems 4 and 5. ■

5.2 Efficiency

To indicate the dependence of γ^i and Γ on δ, we write γ^i_δ and Γ_δ. Observe that Γ_δ is bounded and increasing in δ. In fact, if $\Gamma_\delta(d) > 0$ for some d and some δ, then $\Gamma_{\delta'}(d) \geq d$ for δ' sufficiently large, implying that d^* is bounded away from zero for discount factors close enough to one. If this is the case, then any action profile can be supported in a single period if the players are sufficiently patient. This argument yields a necessary and sufficient condition for patient players to attain payoffs on the efficient frontier of the payoff set.

Theorem 6. *If $\Gamma(\infty) > 0$ then $d^* > 0$ and $V(S_{\text{CE}})$ is a subset of the efficient frontier for δ sufficiently close to 1. If $\Gamma(\infty) = 0$ then $V(S_{\text{CE}})$ attains the same welfare level as the welfare-maximizing Nash equilibrium of the stage game.*

Proofs for this subsection are in the Supplemental Appendix.

Example 1 (The Prisoners' Dilemma). Theorem 6 is illustrated by the Prisoners' Dilemma:

	C	D
C	1, 1	−r, x
D	x, −r	0, 0

To apply Theorem 6, we observe that

$$\gamma^2(\infty) = \max\{\pi_2 - \pi_1 - (x-1),\ \pi_2 x - \pi_2 r,\ -\pi_1 x - \pi_2 r,\ 0\}, \tag{15}$$

$$\gamma^1(\infty) = \max\{\pi_1 - \pi_2 - (x-1),\ -\pi_2 x - \pi_1 r,\ \pi_1 x - \pi_1 r,\ 0\}, \tag{16}$$

where the elements in each maximand are the values at CC, DC, CD, and DD, in this 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 $\gamma^2(\infty)+\gamma^1(\infty) > 0$ if and only if either $x > r$ or $|\pi_2 - \pi_1| > x-1$, which by Theorem 6 is necessary and sufficient for CC to be supported in contractual equilibrium for sufficiently high $\delta < 1$. If, on the other hand, $x < r$ and $|\pi_2 - \pi_1| < x - 1$, then $W^* = \{(0,0)\}$ regardless of δ.

To get a feel for the construction of the set W^*, consider Figure 1. Suppose that $\pi_1 = \pi_2$ and $x > r$. Consider player 1's favorite point z^2. It will be the case that z^2 is achieved with reference to the disagreement point $\underline{w}^2$ that is furthest in the direction $(\pi_2, -\pi_1) = \left(\frac{1}{2}, -\frac{1}{2}\right)$. This disagreement point, in turn, will be the weighted average of $u(\mathrm{DC}) = (x, -r)$ and a point $v' \in W^*$. 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'. To discourage player 2's deviation, it must be that $\delta(v'_2 - z^2_2) \geq (1-\delta) r$. The construction is pictured in Figure 1. Note that $\underline{w}^2$ is the disagreement point and it leads to z^2 as the solution to the bargaining problem.

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

Theorem 7. *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*

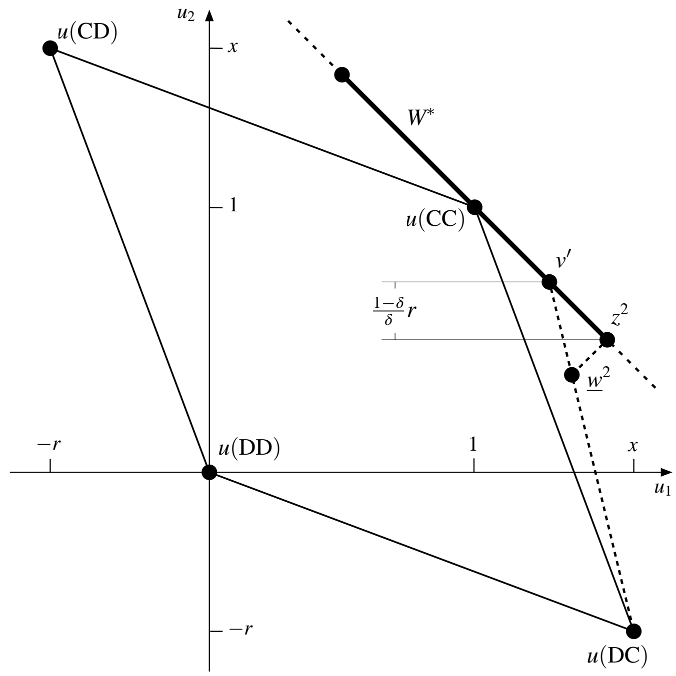

FIGURE 1. THE PRISONERS' DILEMMA GAME. The endpoint z^2 is attained by playing CC in the stage game and using a transfer to split the surplus relative to the disagreement point $\underline{w}^2$. The disagreement point $\underline{w}^2$, in turn, is attained by playing DC in the stage game and continuing with promised utility v' if no deviation occurs, and promised utility z^2 if player 2 deviates. Maximizing z_1^2 subject to enforcing DC under disagreement requires $v_2' - z_2^2 = \frac{1-\delta}{\delta}r$.

yields a BSG 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 i.*
- *Disagreement: In state i, play a^i. If player $j \neq i$ deviates unilaterally, go to state j; otherwise stay in state i.*

The continuation values in the two states are the endpoints of a BSG set that is contained in $V(S_{\text{CE}})$.

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 state $-i$. 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 π_2/π_1-sloped line from their stage game payoff under disagreement.

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 ζ^i be the pure action minimax payoff profile for player i in the stage game. Suppose that $(\pi_2, -\pi_1) \cdot (\zeta^2 - \zeta^1) > 0$. Then contractual equilibrium attains efficiency 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 8. *Suppose that $u : \mathcal{A} \to \mathbb{R}$ is uniformly bounded, and for each player i there exists an open interval $(\underline{a}_i, \overline{a}_i) \subset \mathcal{A}_i$ such that (i) u is twice differentiable on $(\underline{a}_1, \overline{a}_1) \times (\underline{a}_2, \overline{a}_2)$; (ii) each player i's best response function BR_i is differentiable on $(\underline{a}_1, \overline{a}_1) \times (\underline{a}_2, \overline{a}_2)$; and (iii) there exists a stage game Nash equilibrium $a^{\mathrm{NE}} \in (\underline{a}_1, \overline{a}_1) \times (\underline{a}_2, \overline{a}_2)$. If, for both i, $dBR_i(a_{-i})/da_{-i} < 1$ and $du_i(a)/da_{-i} < 0$ at a^{NE}, then contractual equilibrium attains efficiency 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 and a uniformly bounded profit function, its best response functions have negative slope, and each firm's payoff is decreasing in the other's quantity. That is, it satisfies the conditions of Theorem 8, and therefore has an efficient contractual equilibrium if the firms are sufficiently patient.

5.3 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 span of $V(S_{\mathrm{CE}})$, and hence the welfare level attainable in contractual equilibrium. In fact, in general the highest welfare level is attained when bargaining power is extremely unequal. This feature is on prominent display in the applications described in Section 7.

Theorem 9. *Given $\mathcal{A}$, δ, and u,* $\mathrm{level}(V(S_{\mathrm{CE}}))$ *is maximized at either $\pi = (0, 1)$ or $\pi = (1, 0)$.*

The proof is in the Supplemental Appendix. Intuitively, any BSG set is the projection of the relevant disagreement points onto the frontier in the direction of the bargaining shares. Since equal bargaining shares form a vector perpendicular to the frontier, they minimize the projected distance between the disagreement points. This property is evident in the Prisoners' Dilemma (Example 1). Recall that in the Prisoners' Dilemma a sufficient condition for efficiency at high discount factors is $|\pi_2 - \pi_1| > x - 1$, which is most relaxed when $\pi_1 = 1$ or $\pi_2 = 1$. The following example provides a more direct illustration.

Example 2 (The Prisoners' Double Dilemma). The stage game of the Prisoners' Double Dilemma,

	C	D	E
C	$1, 1$	$-3, 2$	$-3, 2$
D	$2, -3$	$0.2, 0.2$	$0, 0$
E	$2, -3$	$0, 0$	$0, 0$

is similar to the Prisoners' Dilemma. Since $x - r = 2 - 3 < 0$, 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 a BSG set with a payoff span of 0, and thus efficiency is not attainable. However, whenever $\pi \neq \left(\frac{1}{2}, \frac{1}{2}\right)$ Figure 2 illustrates that playing DD and EE under disagreement supports a BSG set with a strictly positive payoff span, and thus efficiency is attainable for sufficiently high δ.

6 Generalization

In this section, we extend our analysis to a general model with more than two players, imperfect public monitoring, and heterogeneous discount factors.[15] For simplicity, we do not specify the bargaining process, and instead assume that there exist unique bargaining weights that summarize every backward induction solution of the bargaining phase under the axioms. (It would be cumbersome, but straightforward, to construct such a bargaining process using a generalized alternating-offer protocol.) A simplified description of a game in this class is a tuple $\langle n, \mathcal{A}, \Theta, f, u, \vec{\delta}, \pi\rangle$, with:

- A finite number of players n;
- A stage game, featuring a set of action profiles $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$, a set of public signals Θ, a signal distribution function $f\colon \mathcal{A} \to \Delta\Theta$, and payoff functions $\left(u_i : \mathcal{A}_i \times \Theta \to \mathbb{R}\right)_{i=1}^{n}$;

[15]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).

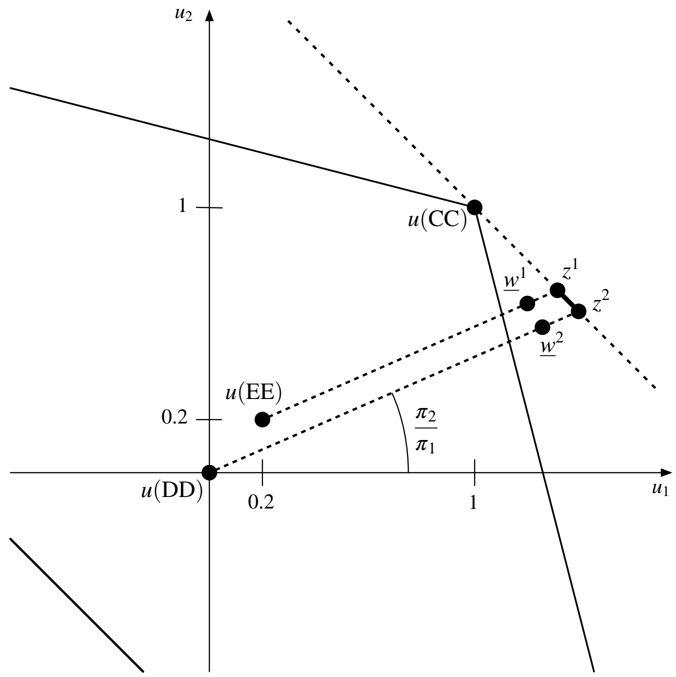

FIGURE 2. THE PRISONERS' DOUBLE DILEMMA GAME. The endpoint z^2 is attained by playing CC in the stage game and using a transfer to split the surplus relative to the disagreement point $\underline{w}^2$. The disagreement points $\underline{w}^1$ and $\underline{w}^2$ are attained by playing stage game Nash equilibria EE and DD, respectively, and continuing with promised utility z^1 and z^2, respectively, regardless of whether either player deviated.

- A vector of discount factors $\vec{\delta} \in [0, 1)^n$, where δ_i denotes player i's discount factor;
- 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 their actions, yielding action profile $a \in \mathcal{A}$. Public signal $\theta \in \Theta$ is then realized according to the probability measure $f(\cdot|a)$. The players publicly observe θ but they do not observe each other's actions. We write $u(a, \theta) \equiv \big(u_i(a_i, \theta)\big)_{i=1}^n$, and extend f and u to the space of mixed actions.

Since the players' common 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 the action phase, the players condition on the history through only the public signals realized in previous periods. That is, 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 $\tilde{w}$ from the negotiation phase must now incorporate (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 the component-by-component multiplication operator $*$ as follows:

$$x * x' \equiv (x_1 x'_1, x_2 x'_2, \ldots, x_n x'_n). \tag{17}$$

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 $\tilde{w}$ is constructed as follows:

$$\tilde{w} = m + \int_{\theta \in \Theta} \big(u(\alpha, \theta) + \vec{\delta} * \tilde{g}(\theta)\big)\, df(\theta|\alpha), \tag{18}$$

where $m \in \mathbb{R}^n_0 \equiv \{m' \in \mathbb{R}^n \mid \sum_{i=1}^n m'_i = 0\}$ is the transfer, $\alpha \in \Delta^{\mathrm{U}}\mathcal{A}$ is the mixed action profile, $\tilde{g}(\theta)$ is the continuation value from the start of the next period following public signal θ, and $f(\theta|\alpha)$ is the probability measure on Θ that arises from α. A disagreement value is given by:

$$\underline{\tilde{w}} = \int_{\theta \in \Theta} \big(u(\alpha, \theta) + \vec{\delta} * \tilde{g}(\theta)\big)\, df(\theta|\alpha). \tag{19}$$

We say that $\tilde{g} : \Theta \to \tilde{W}$ enforces α if α is a Nash equilibrium of the game with action-profile space $\mathcal{A}$ and payoffs given by

$$\int_{\theta \in \Theta} \big(u(a, \theta) + \vec{\delta} * \tilde{g}(\theta)\big)\, df(\theta|a). \tag{20}$$

Operators D, C, and B are revised as follows:

$$D(\tilde{W}) \equiv \left\{ \underline{\tilde{w}} \in \mathbb{R}^n \middle|\, \exists\, \tilde{g} : \Theta \to \tilde{W} \text{ and } \alpha \in \Delta^{\mathrm{U}}\mathcal{A} \text{ s.t. } \tilde{g} \text{ enforces } \alpha \text{ and Eq. 19 holds} \right\}, \tag{21}$$

$$C(\tilde{W}) \equiv \left\{ \tilde{w} \in \mathbb{R}^n \middle|\, \exists\, m \in \mathbb{R}^2_0 \text{ and } \tilde{w}' \in D(W) \text{ s.t. } \tilde{w} = m + \tilde{w}' \text{ and } \tilde{w} \in \overline{D}(\tilde{W}) \right\}, \tag{22}$$

where $\overline{D}$ is defined as before, and

$$B(\tilde{W}, \underline{\tilde{W}}) \equiv \left\{ \underline{\tilde{w}} + \pi \Big(\max_{\tilde{w}' \in \tilde{W}} \vec{1} \cdot \tilde{w}' - \vec{1} \cdot \underline{\tilde{w}} \Big) \ \Big| \ \underline{\tilde{w}} \in \underline{\tilde{W}} \right\}. \tag{23}$$

As before, $\tilde{W}$ is a BSG set if $\tilde{W} = \operatorname{co} B(C(\tilde{W}), D(\tilde{W}))$; the definition of full Pareto dominance carries over unchanged as well.

With $\mathcal{A}$ and Θ finite, we can guarantee existence.

Theorem 10. *Consider any n-player game in the simplified form $\langle n, \mathcal{A}, \Theta, f, u, \vec{\delta}, \pi \rangle$. If $\mathcal{A}$ and Θ are finite and $\delta_i \in [0, 1)$ for all i, then the game has a unique dominant BSG set $\tilde{W}^*$. Moreover, $\tilde{W}^*$ is a compact hyperpolygon contained in a hyperplane normal to the vector $\vec{1}$.*

The proof, in the Supplemental Appendix, 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 $\mathbb{R}^n$. 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

For the special case of two players, we give a characterization of $\tilde{V}(S_{\text{CE}})$ along the lines of the construction in Section 5.1, allowing for imperfect public monitoring and asymmetric different discount factors $\vec{\delta} = (\delta_1, \delta_2)$. The proof is in the Supplemental Appendix.

Theorem 11. *For a two-player game in the simplified form $\langle 2, \mathcal{A}, \Theta, f, u, \vec{\delta}, \pi \rangle$, if $\mathcal{A}$ and Θ are finite and $\tilde{W}^* = \tilde{V}(S_{\text{CE}})$, then $\tilde{W}^*$ is a compact line segment of slope -1, satisfying:*

1. $\operatorname{span}(\tilde{W}^*)$ *is equal to the maximal fixed point of $\tilde{\Gamma} \equiv \tilde{\gamma}^2 + \tilde{\gamma}^1$, where, for players $i \neq j$ and defining $\psi \equiv \pi_1 \delta_2 + \pi_2 \delta_1$ and $\hat{u}(\alpha) \equiv \int_{\theta \in \Theta} u(\alpha, \theta)\, df(\theta|\alpha)$,*

$$\tilde{\gamma}^j(\tilde{d}) \equiv \max_{\eta, \alpha} \left(\frac{\pi_j}{1-\psi} \hat{u}_i(\alpha) - \frac{\pi_i}{1-\psi} \hat{u}_j(\alpha) + \frac{\psi}{1-\psi} \hat{\eta}(\alpha) \right)$$
$$\text{s.t.} \begin{cases} \eta : \Theta \to [-\tilde{d}, 0], \text{ with } \hat{\eta}(\alpha) \equiv \sum_{\theta \in \Theta} f(\theta|\alpha) \eta(\theta) \\ \alpha \in \Delta^{\text{U}}\mathcal{A} \text{ is a Nash equilibrium of } \big\langle \mathcal{A}_i \times \mathcal{A}_j, (\hat{u}_i, \hat{u}_j) + \vec{\delta} * (\hat{\eta}, -\hat{\eta}) \big\rangle; \end{cases} \tag{24}$$

2. $\text{level}(\tilde{W}^*)$ *is equal to*

$$\frac{1-\psi}{(1-\delta_1)(1-\delta_2)}\begin{pmatrix}\delta_1\tilde{\gamma}^2(\text{span}(\tilde{W}^*))+\delta_2\tilde{\gamma}^1(\text{span}(\tilde{W}^*))\\ -\ \delta_2\,\text{span}(\tilde{W}^*)+\chi(\text{span}(\tilde{W}^*))\end{pmatrix}, \tag{25}$$

where

$$\begin{aligned}\chi(\tilde{d}) = \max_{\eta,\alpha}\ & \hat{u}_1(\alpha)+\hat{u}_2(\alpha)+(\delta_1-\delta_2)\hat{\eta}(\alpha)\\ \text{s.t.}\ & \begin{cases}\eta:\Theta\to[-\tilde{d},0],\ \text{with}\ \hat{\eta}(\alpha)\equiv\sum_{\theta\in\Theta}f(\theta|\alpha)\eta(\theta),\\ \alpha\in\Delta^{\text{U}}\mathcal{A}\ \text{is a Nash equilibrium of}\ \langle\mathcal{A},\hat{u}+\vec{\delta}*(\hat{\eta},-\hat{\eta})\rangle.\end{cases}\end{aligned} \tag{26}$$

3. *The endpoints* $\tilde{W}^*$ *are*

$$\tilde{z}^1 = (-1,1)\,\tilde{\gamma}^1\big(\text{span}(\tilde{W}^*)\big)+\left(\frac{1-\delta_2}{1-\psi},\frac{1-\delta_1}{1-\psi}\right)*\pi\,\text{level}(\tilde{W}^*), \tag{27}$$

$$\tilde{z}^2 = (1,-1)\,\tilde{\gamma}^2\big(\text{span}(\tilde{W}^*)\big)+\left(\frac{1-\delta_2}{1-\psi},\frac{1-\delta_1}{1-\psi}\right)*\pi\,\text{level}(\tilde{W}^*). \tag{28}$$

7 Application to relational contracting

7.1 Example 1: 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 stage game, the agent chooses effort $e \in [0,\bar{e}]$, incurring a cost $c(e)$, and then the principal makes a voluntary payment to the agent. The principal's revenue is a random variable θ, with probability density $f(\cdot|e)$ and full support on a compact interval. The principal does not observe e, but θ is public, and so the principal's voluntary payment can be conditioned on θ. We assume that c is strictly increasing and strictly convex, $c(0)=0$, f has the monotone likelihood property, and $f\big(\theta\big|e=c^{-1}(\cdot)\big)$ is convex. The players engage in bargaining and can make voluntarily transfers before the stage game in each period, 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. What is the maximum effort that can be sustained in equilibrium? The answer depends on the allocation of bargaining power, where $\pi_{\text{A}} \geq 0$ is the

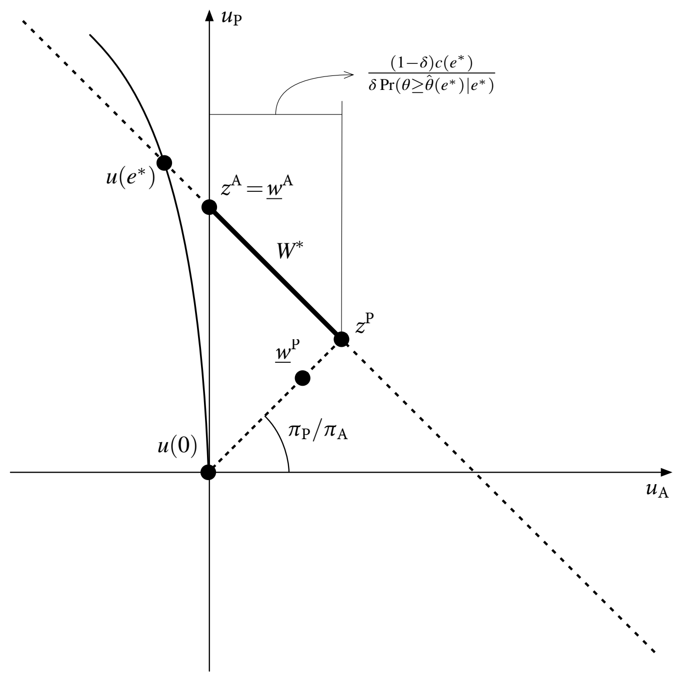

FIGURE 3. THE PRINCIPAL-AGENT GAME. The disagreement point $\underline{w}^P$ is attained by playing stage game Nash equilibrium $e = 0$ and continuing with promised utility z^P regardless of θ. The disagreement point $\underline{w}^A$ is attained by playing $e = e^*$ in the stage game, continuing with promised utility z^P if $\theta \geq \hat{\theta}(e^*)$, where e^* is the maximum effort supported by the difference between z^P_A and z^A_A, and continuing with promised utility z^A otherwise.

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 P rewards the agent and state A punishes the agent. The construction is illustrated in Figure 3.

In state P, 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 P, the players stay in state P the following period regardless of the realized output.[16] Under agreement in state P, the players recognize that their "outside option" is to disagree, implement zero effort, and then return to state P 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 P. Since the agent obtains a π_{A} fraction of the surplus, their agreement utility is $(\pi_{\mathrm{A}}, 1-\pi_{\mathrm{A}}) \cdot (\mathbb{E}(\theta|e^*) - c(e^*))$, 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 A, 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 A for low realizations of θ, and transiting to state P for high realizations of θ.[17] As in Levin (2003), the optimal cutoff for enforcing effort level e^* is the output level $\hat{\theta}(e^*)$ at which $\frac{\partial}{\partial e} f(\theta|e)|_{e=e^*}$, as a function of θ, switches from negative to positive (by the monotone likelihood property, this identifies a unique output level). In fact, this same incentive scheme must be used whenever the agent is to exert effort e^*, i.e., in state P under agreement as well as in state A under both agreement and disagreement.

Since optimal effort is played under disagreement in state A, 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, 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 A must be at least zero. To provide the maximal incentives, in fact his utility in state A should be exactly zero.

Equilibrium path effort e^* is thus a fixed point of the agent's optimization problem:[18]

$$e^* \in \arg\max_{e \in [0, \bar{e}]} \Big(-(1-\delta)\, c(e) + \delta \Pr\big(\theta \geq \hat{\theta}(e^*) \big| e\big) \pi_A \big(\mathbb{E}(\theta|e^*) - c(e^*)\big)\Big) \tag{29}$$

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}, and converges to efficient effort as

[16] Even if the principal could promise a voluntary payment (conditioned on y) in the action phase, she would not do so in state P under disagreement. Since state 1 is already her worst state, she cannot be punished further for reneging on such a promise.

[17] If the principal could promise a voluntary payment for high realizations of θ in the action phase, it would have the same effect as transitioning to state P.

[18] We assume for convenience that $\bar{e} \in \arg\max_{e \in [0, \bar{e}]} \big[\mathbb{E}(\theta|e) - c(e)\big]$, so inefficiently high effort is infeasible.

$\delta \to 1$ and $\pi_A \to 1$.

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.

Levin's "strongly optimal" equilibrium relies 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. In contractual equilibrium, in contrast, no-fault disagreement ensures a well-defined default point for bargaining each period, while internal and Pareto external agreement consistency allow the agents to endogenously select a division of the surplus that accords with their bargaining power.

7.2 Example 2: External enforcement and self-enforcement

To further illustrate the usefulness of contractual equilibrium, this section reconsiders the classic question posed by Baker, Gibbons, and Murphy (1994): Does external enforcement help or harm self-enforcement? Baker, Gibbons, and Murphy focus on "trigger punishments" in which the agent punishes the principal for reneging on a voluntary payment by refusing to accept "implicit" (self-enforced) contracts offered in the future. When "explicit" (externally enforceable) contracts are available, Baker, Gibbons, and Murphy assume that even in a trigger punishment the agent would still accept an externally enforced contract, and therefore the principal offers the externally enforced contract that maximizes her profit. Their key insight is that the threat of terminating the relationship (in the absence of externally enforced contracts) is a more severe punishment than reverting to profit-maximizing externally enforced contracts, and can therefore support higher payoffs in equilibrium.

However, a trigger punishment (either termination or reversion to explicit contracts) is generally not viable under agreement in contractual equilibrium, because the principal and the agent can bargain their way out of it. Furthermore, neither can an explicit "spot market" contract of the sort considered by Baker, Gibbons, and Murphy arise under disagreement in contractual equilibrium, because it requires the agreement of both parties. Therefore an externally enforced contract cannot serve as their "fallback position" unless it is already in place prior to bargaining.

These concerns lead us to a model of external enforcement motivated by the observation that in salaried employment relationships there is typically an externally enforced contract in place at all times—the employment contract. Employment contracts often have long or indefinite horizons, and, though often subject to voluntary termination (i.e., firing by the principal or quitting by the agent), are held in place by the fact that termination is quite costly. To broaden our results, we generalize the model from Section 7.1 by allowing the agent to exert costly "negative effort" that is verifiable to an external enforcer. Positive effort, as in the previous section, is not verifiable to the external enforcer.[19] At first we take the externally enforced contract to be given exogenously, and discuss later how it might be chosen endogenously.

When the agent has some bargaining power, we find that the availability of external enforcement improves the prospects for self-enforcement. In particular, an optimal externally enforced contract induces the agent to exert negative effort, by compensating him for his verifiable effort costs. This externally enforced contract is always undone by self-enforced incentives under agreement; it affects only behavior under disagreement. A surplus-destroying externally enforced contract raises the stakes of bargaining when the agent is being rewarded, enabling the agent to capture more of the surplus. When the agent is being punished, the same behavior is implemented under disagreement and agreement, so there is no surplus to bargain over. Compared to the case without external enforcement, the difference in the agent's payoff under agreement in the two states is greater, supporting higher effort on the equilibrium path.

The contractual equilibrium we describe is similar to the one outlined in Section 7.1, and is illustrated in Figure 4. In state P under disagreement, the principal pays only to the letter of the externally enforced contract, and the agent plays his stage-game best response. In state A under disagreement and in both states under agreement, the agent plays his equilibrium-path effort, and is optimally rewarded by transiting to state P for good outcomes and punished by transiting to state A for bad outcomes.

Suppose that $e \in [\underline{e}, \overline{e}]$, where $\underline{e} < 0$, and the agent's cost of effort $c(e)$ is convex and minimized at $e = 0$. That is, the agent can exert costly effort to reduce the principal's revenue below $\mathbb{E}(\theta | e = 0)$. Given the externally enforced contract, let $\hat{e}$ be the agent's stage-game best response and let $\hat{u}$ be the expected stage game payoff vector, including the externally enforced contractual payments, that results when the agent plays $\hat{e}$. Because the agent can always deviate to obtain $\hat{u}$, his utility in each state must be at least $\hat{u}$. To provide maximal incentives, of course, his utility in state A, $z_{\mathrm{A}}^{\mathrm{A}}$, should be exactly $\hat{u}$, under both agreement and disagreement.

In state P under agreement, the principal and the agent share the surplus according to their

[19] Our conclusion does not depend on the exact nature of the verifiable element on which externally enforced contracts are based. The negative-effort story makes the theme clear.

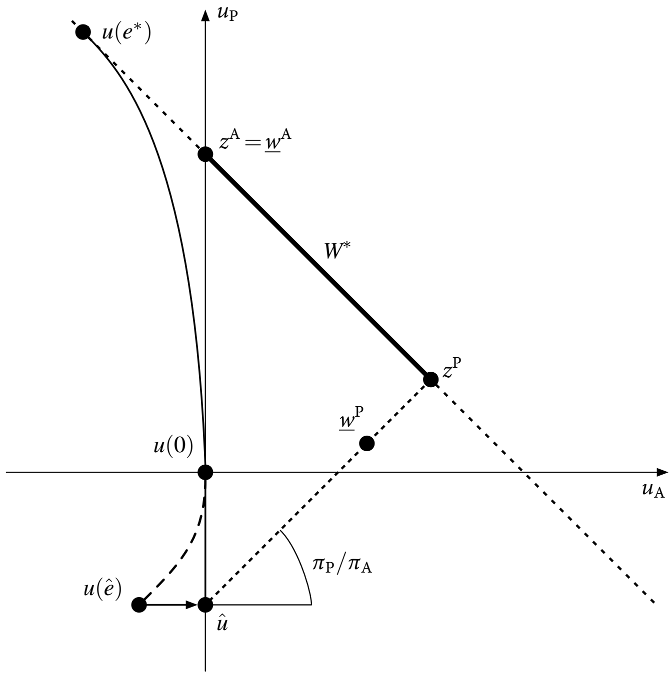

FIGURE 4. THE EXPLICIT CONTRACT GAME. The disagreement point $\underline{w}^{\mathrm{P}}$ is attained by playing stage game Nash equilibrium $e = \hat{e}$ under the externally enforced contract, and then continuing with promised utility z^{P} regardless of θ. The disagreement point $\underline{w}^{\mathrm{A}}$ is attained by playing $e = e^*$ in the stage game, continuing with promised utility z^{P} if $\theta \geq \hat{\theta}(e^*)$, where e^* is the maximum effort supported by the difference between $z^{\mathrm{P}}_{\mathrm{A}}$ and $z^{\mathrm{A}}_{\mathrm{A}}$, and continuing with promised utility z^{A} otherwise. Compared to the Principal-Agent game in Figure 3, the payoff span of W^* is larger, and therefore higher effort is attained on the equilibrium path.

bargaining shares, so the agent's utility is

$$\begin{aligned} z^{\mathrm{P}}_{\mathrm{A}} &= (1-\delta)\hat{u}_{\mathrm{A}} + \delta z^{\mathrm{P}}_{\mathrm{A}} + \pi_{\mathrm{A}}\big(\mathbb{E}(\theta|e^*) - c(e^*) - \big((1-\delta)(\hat{u}_{\mathrm{A}} + \hat{u}_{\mathrm{P}}) + \delta(z^{\mathrm{P}}_{\mathrm{A}} + z^{\mathrm{P}}_{\mathrm{P}})\big)\big) \\ &= \hat{u}_{\mathrm{A}} + \pi_{\mathrm{A}}\big(\mathbb{E}(\theta|e^*) - c(e^*) - (\hat{u}_{\mathrm{A}} + \hat{u}_{\mathrm{P}})\big) \end{aligned} \tag{30}$$

The strength of the agent's incentives are measured by the difference between his payoffs under agreement in state P and state A; i.e., $z^{\mathrm{P}}_{\mathrm{A}} - z^{\mathrm{A}}_{\mathrm{A}} = \pi_{\mathrm{A}}\big(\mathbb{E}(\theta|e^*) - c(e^*) - (\hat{u}_{\mathrm{A}} + \hat{u}_{\mathrm{P}})\big)$. Therefore his incentives are strictly decreasing in $\hat{u}_{\mathrm{A}} + \hat{u}_{\mathrm{P}}$. Intuitively, when joint payoffs are lower under disagreement in state P, there is more to be gained from agreeing, and the agent gets a π_{A} share of the gains.

So far we have assumed that the externally enforced contract is exogenous. But suppose the externally enforced contract is jointly chosen by the principal and the agent in their first-period bargaining phase. Since their joint payoff in equilibrium is strictly decreasing in their joint payoff under disagreement in state P, they want their externally enforced contract to destroy enough surplus to support unconstrained efficient effort on the equilibrium path. In a broader environment, surplus destruction could be constrained by the ability of the principal to fire the agent at some cost and the agent to quit at some cost. If there are tight constraints on surplus destruction, then a straight salary without externally enforced incentives can be optimal.

8 Conclusion

This paper introduces *contractual equilibrium*, a new approach to equilibrium selection in repeated games based on communication with endogenous meaning. Contractual equilibrium allows disagreement play to arise endogenously, in a way consistent with a well-defined extensive form. Contractual equilibrium exists, yields a unique welfare level, is tractable in applications, and extends easily to games with imperfect monitoring and heterogeneous discount factors.

Bargaining power in contractual equilibrium arises from the bargaining protocol, while the surplus available to bargain over arises from the theory of disagreement embodied in the no-fault disagreement axiom. No-fault disagreement posits that disagreement play should not depend on the manner of disagreement, although it may depend on the history of actions and agreements.

The results presented in Section 7 show that contractual equilibrium yields new insights even when applied to classic models from the literature. We hope contractual equilibrium will be viewed as a kind of platform layer, providing tools that simplify and accelerate the study of relational contracts. Without the need to generate new assumptions about bargaining and disagreement for each new application, it will be easier for relational contract models to incorporate more realistic details.

In conjunction with promoting the analysis of relational contract applications, the modeling framework developed here may be useful for interpreting existing models in the literature in terms of implied notions of disagreement. Renegotiation proofness, for instance, can be viewed as assuming that play under disagreement is identical to that under agreement. Some approaches in the relational contracts literature assume that disagreement yields separation, a spot contract, or a stage-game equilibrium. At the other end of the spectrum, disagreement play could depend on anything, including the manner of disagreement, but we show that a substantive theory of disagreement is needed to refine the set of subgame perfect equilibrium payoffs.

The underlying principles of contractual equilibrium are not restricted to repeated games with transfers. 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 a dynamic setting, since the actions taken under disagreement can change the set of feasible continuation values.[20] Furthermore, one can apply the same definitions to games without a transfer phase, although a modified version of no-fault disagreement would be required. In general, bargaining power will arise as long as the theory of disagreement embodies some restriction on play. Finally, it would be interesting to consider how disagreement play might depend on which coalitions manage to cohere when the grand coalition fails to agree. We view all of these extensions as promising areas for further research.

Since the backward induction solution of the generalized alternating-offer bargaining game under our axioms is expressed in terms of bargaining power, it yields the same outcome as the generalized Nash bargaining solution. Therefore the same conclusions would arise from a hybrid model in which the noncooperative bargaining and transfer phases are replaced with a cooperative Nash bargaining phase. There is a tradition of incorporating cooperative bargaining into otherwise noncooperative games in order to model bargaining power. One example is the macroeconomic literature on labor market search frictions, where employment contracts are typically reached through Nash bargaining (see footnote 7). Another leading example is the hold-up 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 trade takes place.[21] Watson (2006) formalizes a general class of finite-horizon hybrid games that nests the earlier literature, and defines a *contractual equilibrium* as a strategy profile that satisfies a cooperative bargaining rule at each cooperative node and subgame perfection at every noncooperative node. Our representation theorem can be interpreted as stating an equivalence between fully noncooperative repeated game equilibria with meaningful bargaining, and contractual equilibria in a corresponding hybrid game with cooperative bargaining. However, one need not take this interpretation to make use of our results.

[20]The legal environment—the set of enforceable contracts and the enforcement technology—can introduce dynamics if long-term contracts are enforceable. Then the players' signatures on the contract, though they have 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.

[21]In addition, Schmidt and Schnitzer (1995) and Ramey and Watson (2002) employ a hybrid approach to relational contracts. Several papers examine the actions that take place leading up to a agreement, where the agreement yields a contract that terminates the game. Brandenburger and Stuart (2007) use cooperative bargaining, Busch and Wen (1995) use noncooperative bargaining, and Jackson and Palfrey (2001) use noncooperative implementation.

A Proofs

A.1 Agreement consistency has no bite

We expand Theorem 1 by invoking a stronger consistency axiom while proving the same result. Since the players recognize that any payoff in $C(\text{co } V(S_{\text{IAC}}))$ is attainable from the action phase, suppose that any proposal to play as if switching to a different equilibrium in S_{IAC} must be interpreted meaningfully.

Axiom SEAC (Strong external agreement consistency)**.** For every history to the transfer phase under agreement, if the agreement takes the form $w \in C(\text{co } V(S_{\text{IAC}}))$ then continuation play yields the value w.

Let S_{SEAC} be the subset of S_{SPE} that satisfies SEAC, and note that $S_{\text{SEAC}} \subseteq S_{\text{IAC}}$. SEAC imposes a strong requirement on continuation play following a deviant agreement, because it requires the players to follow through even on a deviant agreement that is Pareto dominated by what they should have agreed on. However, IAC and SEAC do not constrain the set of attainable payoffs.

Theorem 1′ (IAC and SEAC have no bite)**.** $V(S_{\text{IAC}}) = V(S_{\text{SEAC}}) = V(S_{\text{SPE}})$.

Proof of Theorem 1′ First, since bargaining is cheap talk and there is public correlation, $V(S_{\text{NB}}) = V(S_{\text{SPE}})$. Thus, our task is to prove $V(S_{\text{NB}}) \subset V(S_{\text{SEAC}})$. We proceed by demonstrating how, for a given period, values in $C(V(S_{\text{NB}}))$ can be supported from the bargaining phase in a way that is consistent with SEAC. We then construct equilibrium strategies recursively.

Consider any $w \in C(V(S_{\text{NB}}))$. By the definition of C, there is a specification of behavior for the transfer phase and action phase that (i) yields value w from the transfer phase, (ii) utilizes continuation values from $V(S_{\text{NB}})$ for the following period, and (iii) is sequentially rational. We refer below to this specification as "rational behavior supporting w from the transfer phase."

For any fixed value in $C(V(S_{\text{NB}}))$, we next describe behavior in the bargaining phase that yields this value. We specify the behavior recursively by round of the bargaining phase, starting with the first round. Take as given a round of the bargaining phase and any $w \in C(V(S_{\text{NB}}))$. Call w the "desired value." Let i denote the proposer and let j denote the responder. The following provisions give the prescribed behavior for this round of bargaining.

1. Prescribe that player i proposes w and player i responds "yes." If the players behave in this way, then prescribe rational behavior supporting w from the transfer phase.
2. Suppose player i proposes w as prescribed, but player j deviates by responding "no." If bargaining then breaks down, prescribe rational behavior supporting w from the transfer phase. Otherwise, let w again be the desired value for the next round of bargaining.
3. Suppose player i deviates by proposing some $w' \notin C(V(S_{\text{NB}}))$. If bargaining terminates after player j's response, either because player j responds "yes" or because bargaining breaks down, then prescribe rational behavior supporting w from the transfer phase. If bargaining does not terminate, then let w be the desired value for the next round of bargaining.

4. Suppose player i deviates by proposing some $w' \in C(V(S_{\mathrm{NB}}))$. Then select $w'' \in C(V(S_{\mathrm{NB}}))$ such that $w''_i \leq w_i$ and $w''_j \geq w'_j$. (By the characterization of C given by Eq. 5, w'' can always be attained by augmenting either w or w' with a transfer from player i to player j.) Prescribe that player j responds "no" to the offer w'. If player j instead deviates to respond "yes," then prescribe rational behavior supporting w' from the transfer phase. If player j responds "no" and bargaining breaks down, then prescribe rational behavior supporting w'' from the transfer phase. If player j responds "no" and bargaining continues, then set the desired value to w'' for the next round of bargaining.

In some of the cases just described, bargaining continues to another round. Behavior for the next round is specified in the same fashion, with the new desired value. In this way, we recursively construct a complete specification of behavior in the bargaining phase, as well as for the transfer phase and the action phase. There is an implied mapping from the sequence of actions in the period to continuation values from the start of the next period.

The specified behavior is sequentially rational. For instance, by proposing the desired value w, player i expects to get w_i. If he deviates then he gets either w_i or some $w''_i \leq w_i$, so he prefers to propose w. Note that in the event that player i proposes some $w' \in C(V(S_{\mathrm{NB}}))$ (even a value satisfying $w'_j > w_j$), player j optimally says "no" because this leads to the value w'' which player j prefers to w'.

The specified behavior is also consistent with SEAC. To see this, first note that the presence of the public randomization phase implies $\operatorname{co} V(S_{\mathrm{NB}}) = V(S_{\mathrm{NB}})$. In addition, since $S_{\mathrm{IAC}} \subset S_{\mathrm{NB}}$, we have that $\operatorname{co} V(S_{\mathrm{IAC}}) \subset \operatorname{co} V(S_{\mathrm{NB}})$. Putting these together, and using the fact that C is monotone, we conclude that $C(\operatorname{co} V(S_{\mathrm{IAC}})) \subset C(V(S_{\mathrm{NB}}))$. Thus, any agreement on a value in $C(\operatorname{co} V(S_{\mathrm{IAC}}))$ is covered by item 1 or item 4, which specify behavior that attains the agreed value.

The construction above shows that all values in $C(V(S_{\mathrm{NB}}))$ can be supported from the bargaining phase in a way consistent with SEAC, using continuation values in $V(S_{\mathrm{NB}})$ from the following period. Thus, SEAC does not impose a constraint on the recursive construction that characterizes SPE continuation values, and $V(S_{\mathrm{SEAC}}) = V(S_{\mathrm{SPE}})$. For each point $v \in V(S_{\mathrm{SPE}})$, an equilibrium strategy profile that attains value v and satisfies SEAC can be constructed recursively by following the steps outlined above for behavior within a period, and then over all periods (by tracking the continuation values required). ■

A.2 No-fault disagreement has little bite

Theorem 2 focuses on the Pareto frontier of $V(S^{\mathrm{ps}}_{\mathrm{NFD}})$, but here we prove a stronger result that addresses the interior of $V(S^{\mathrm{ps}}_{\mathrm{NFD}})$ as well. Let $S^{\mathrm{nt}}_{\mathrm{SPE}}$ be the set of subgame perfect equilibria that specify zero transfers after all histories.

Theorem 2′ (NFD has little bite)**.** $P(V(S^{\mathrm{ps}}_{\mathrm{SPE}})) = P(V(S^{\mathrm{ps}}_{\mathrm{NFD}})) \subseteq V(S_{\mathrm{NFD}})$ and $V(S^{\mathrm{nt}}_{\mathrm{SPE}}) \subseteq V(S_{\mathrm{NFD}})$.

Proof of Theorem 2′ Standard arguments establish that the sets of interest are closed, so we can focus on closed sets of continuation values W. Because of public randomization, we can also assume that W is convex.

Operators C and $\hat{C}$ share a useful property: For any W, the Pareto frontiers of $C(W)$ and $\hat{C}(W)$ are subsets of a line segment with slope -1. In fact, they have the same extremal values. To see this, suppose that w is the point that minimizes player 1's value among points in $P(C(W))$, so that $w_1 = \min \overline{D}_1(W)$. There is then a point $\underline{w} \in D(W)$ such that $\underline{w}_1 = w_1$. Furthermore, it must be that $\underline{w}_2 \leq w_2$, for otherwise w would not be on the Pareto frontier of $C(W)$. This implies that $w \in \hat{C}(W)$ as well. The same logic applies to the other extremal point. Since $\hat{C}(W) \subset C(W)$, the extrema coincide. We thus have that $P(\text{co}\, C(W)) = P(\text{co}\, \hat{C}(W))$. It is also clear that each player's greatest and least values in $\text{co}\, C(W)$ and $\text{co}\, \hat{C}(W)$ are found on the Pareto frontiers of these sets.

With these facts in mind, we can restrict attention to sets W that have the same properties, since our objects of interest are the fixed points of $\text{co}\, C(W)$ and $\text{co}\, \hat{C}(\cdot)$. With pure strategies, for any history to the action phase the players will be coordinating on a pure action profile. It is sufficient to punish a unilateral deviation by selecting a continuation value that is worst for the deviating player. Furthermore, to find the maximal joint continuation value from the action phase, on the equilibrium path the continuation value from the next period will be on the Pareto frontier of the set of continuation values. Thus, if W is the set of continuation values from the next period then, in reference to Eq. 3, it is sufficient to have g map to three points on the frontier line segment, including the two endpoints. An implication of these observations is that $P(\text{co}\, C(W))$ and $P(\text{co}\, \hat{C}(W))$ depend on only the Pareto frontier of W. Thus, $P(V(S^{\text{ps}}_{\text{SPE}}))$ is itself a fixed point of $\text{co}\, C(\cdot)$ and also of $\text{co}\, \hat{C}(\cdot)$, which implies the first claim of the theorem.

For the second claim, observe that any subgame perfect equilibrium value obtained without transfers can be obtained under NFD if the players ignore the bargaining phase entirely. ■

A.3 Representation

Theorem 3 is proven by the following four lemmas.

Lemma 6. *Consider an isolated noncooperative bargaining game with recognition process ρ, breakdown process β, default payoff $\underline{w}$, and bargaining set $\{w \in \mathbb{R}^2 \mid w_1 + w_2 \leq K\}$ for some $k \in \mathbb{R}$ satisfying $K \geq \underline{w}_1 + \underline{w}_2$. In any subgame perfect equilibrium, with probability 1 each player i earns a payoff of $\underline{w}_i + \pi_i(K - \underline{w}_1 - \underline{w}_2)$, where $\pi_i \equiv \sum_{\ell=1}^{\infty} \left(\rho_{i,\ell}\beta_\ell \prod_{k=1}^{\ell-1}(1-\beta_k)\right)$.*

Proof. This follows from generalizing the recursive methods of Shaked and Sutton (1984) (see also Binmore, Rubinstein, and Wolinsky (1986)), using our assumption that $\prod_{\ell=1}^{\infty}(1-\beta_\ell) = 0$. □

For the rest of the proof, fix $\pi_i = \sum_{\ell=1}^{\infty} \rho_{i,\ell}\beta_\ell \prod_{k=1}^{\ell-1}(1-\beta_k)$. Given a BSG set, we next construct a subgame perfect equilibrium that attains a value in this set while satisfying IAC and NFD.

Lemma 7. *Let W be a BSG set, and select any payoff vector $v \in W$. Inputting (v, W) into the recursive construction algorithm in Figure 5 yields a subgame perfect equilibrium $s \in S_{IAC} \cap S_{NFD}$ that attains an expected payoff of v.*

Input: W a BSG set and $v \in W$.
Output: $s \in S_{\mathrm{IAC}} \cap S_{\mathrm{NFD}}$ such that $v \in V(s)$.

INITIALIZATION:
let $\hat{v}(\varnothing) = v$;
let $\Omega = B(C(W), D(W))$;

for $t = 1$ **to** ∞ **do**
  foreach $h \in \left(\Omega \times \mathcal{B} \times \mathbb{R}_+ \times \mathbb{R}_+ \times \mathcal{A}\right)^{t-1}$ **do**
    PUBLIC RANDOMIZATION:
    let $\mu \in \Delta\Omega$ *solve* $\int_{w \in \Omega} w \, d\mu = \hat{v}(h)$;
    play *public randomization, resulting in realization* w;

    BARGAINING PHASE:
    let $\underline{w} \in D(W)$ *solve* $w \in B\big(C(W), \{\underline{w}\}\big)$;
    play *a subgame perfect equilibrium in an isolated noncooperative bargaining game with recognition process* ρ*, breakdown process* β*, bargaining set* $C(W)$*, and default payoff* $\underline{w}$;

    TRANSFER PHASE:
    foreach $b \in \mathcal{B}$ **do**
      if *b ends in agreement on payoff* $w' \in C(W)$ **then**
        let $(m^*, w^*) \in \mathbb{R}^2_0 \times D(W)$ *solve* $(1-\delta)m^* + w^* = w'$;
      else
        let $m^* = (0,0)$ *and* $w^* = \underline{w}$;
      let $\hat{m}^* = (-\min\{0, m^*_i\})_{i=1}^2$;
      play *transfer* $\hat{m}^*$;

      ACTION PHASE:
      foreach $\hat{m} \in \mathbb{R}^2_+$ **do**
        if *b ends in agreement on* $w' \in C(W)$ **then**
          foreach $i \in \{1,2\}$ **do**
            if $\hat{m}_i < \hat{m}^*_i$ *and* $\hat{m}_{-i} = \hat{m}^*_{-i}$ **then**
              let $w^* \in D(W)$ *be redefined to solve* $w^*_i \leq w'_i$;
        let $\alpha^* \in \Delta^{\mathrm{U}}\mathcal{A}$ *and* $g^* : \mathcal{A} \to W$ *solve* $(1-\delta)u(\alpha^*) + \delta g^*(\alpha^*) = w^*$ *subject to* α^* *is a Nash equilibrium in the game* $\langle \mathcal{A}, (1-\delta)u + \delta g^* \rangle$;
        play *mixed action* α^*;
        foreach $a \in \mathcal{A}$ **do**
          let $\hat{v}(h, \mu, b, \hat{m}, a) = g^*(a)$;

FIGURE 5. RECURSIVE CONSTRUCTION ALGORITHM

Proof. Note that the algorithm recursively constructs function $\hat{v} : \mathcal{H} \to W$ along with the strategy profile s, starting with the null history, then one-period histories, and so on. In line 5 of Figure 5, μ exists because $v \in \text{co}\, B\big(C(W), D(W)\big)$, which follows from W being a BSG set. In line 7, $\underline{w}$ exists by definition of B. In line 8, bargaining ends in agreement on payoff w, by construction of w and Lemma 6. In line 11, m^* and w^* exist by definition of C. In line 19, if the test condition is satisfied for $i = j$ then it is violated for $i = -j$. In line 20, w^* exists because $w' \in C(W)$ satisfies $w' \in \overline{D}(W)$ by definition of C. In line 21, since $w^* \in D(W)$, α^* and g^* exist by definition of D. Therefore the algorithm is well defined.

Collecting the output of all the **play** statements yields a well-defined strategy profile (where line 6 is a move of Nature), since every history is accounted for. By construction, this equilibrium is subgame perfect, since it calls for sequentially rational actions in every bargaining phase (by Lemma 6), sequentially rational transfers in every transfer phase (by lines 10–13 and 16–20), and sequentially rational mixed actions in every action phase (by line 21).

Axiom NFD is satisfied by line 13, which sets all disagreements to yield the same disagreement payoff $\underline{w}$. Axiom IAC is satisfied by lines 10–11 and 17–20, which ensure that any agreement payoff $w' \in C(W)$ with $w' \in \overline{D}(W)$ is implemented. Therefore if W is a BSG set and $v \in W$, the output is a strategy profile $s \in S_{\text{IAC}} \cap S_{\text{NFD}}$. The equilibrium payoff is $v \in V(s)$ by construction. □

Lemma 8. *If $s \in S_{IAC} \cap S_{NFD}$ then* co $V(s)$ *is a BSG set.*

Proof. Since $s \in S_{\text{NFD}}$, every bargaining phase has a well-defined disagreement payoff. Since $s \in S_{\text{IAC}}$, every proposal $w \in C(\text{co}\, V(s))$, if accepted, leads to continuation payoff w. Standard recursive techniques for characterizing equilibrium continuation values in bargaining games (for instance, relating the proposer's supremum values in one round to the players' supremum continuation values from the next round) establish that the bargaining phase has a unique equilibrium outcome and that it has the same representation as in Lemma 6. Further, co $V(s)$ must have a closed Pareto boundary; otherwise existence is violated in the bargaining phase. By sequential rationality, the disagreement payoff $\underline{w}$ after any history satisfies $\underline{w} \in D(V(s)) \subseteq D(\text{co}\, V(s))$. Thus the equilibrium payoff from the bargaining phase is $w \in B\big(C(\text{co}\, V(s)), D(\text{co}\, V(s))\big)$, so every continuation value from the start of a period satisfies $v \in \text{co}\, B\big(C(\text{co}\, V(s)), D(\text{co}\, V(s))\big)$. Hence co $V(s)$ is a BSG set. □

For the following two lemmas, let W^* be the dominant BSG set.

Lemma 9. *Inputting W^* and any $v \in W^*$ into the recursive construction algorithm of Figure 5 yields a strategy profile $s^* \in S_{PEAC}$ such that $v \in V(s^*)$.*

Proof. By definition, W^* fully Pareto dominates every BSG set. By Lemma 8, therefore W^* fully Pareto dominates $V(s)$ for every equilibrium $s \in S_{\text{IAC}} \cap S_{\text{NFD}}$. Since W^* is also a BSG set, every $v \in W^*$ is available in each bargaining phase. Since the recursive construction algorithm always produces an equilibrium $s \in S_{\text{IAC}} \cap S_{\text{NFD}}$, with input W^* it produces an equilibrium $s^* \in S_{\text{PEAC}}$. □

Lemma 10. *If $s^* \in S_{PEAC}$, then $V(s^*) \subseteq W^*$.*

Proof. By Lemma 8, co $V(s^*)$ is a BSG set. By PEAC and Lemma 9, every payoff in W^* is in the bargaining set of every bargaining phase. By sequential rationality, therefore, no $v \in V(s^*)$ can be strictly Pareto dominated by any $v' \in W^*$. This implies that if W^* fully Pareto dominates $V(s^*)$ then it must be that $V(s^*) \subseteq W^*$. Finally, since $W^* = \text{co}\, W^*$, co $V(s^*) \subseteq W^*$. □

A.4 Existence and construction

Proof of Lemma 1. By construction of B, for any set W the elements of $B(C(W), D(W))$ are on a line of slope -1, so co $B(C(W), D(W))$ is a line segment of slope -1. $W \subseteq B(C(W), D(W))$ implies that $W \subseteq V(S_{\text{SPE}})$, wherein each $v \in V(S_{\text{SPE}})$ is bounded above by $v_1 + v_2 \leq \max_{\alpha \in \Delta^{\text{U}}\mathcal{A}}(u_1(\alpha) + u_2(\alpha))$, and below by $v_i \geq \min_{\alpha_{-i} \in \Delta\mathcal{A}_{-i}} \max_{\alpha_i \in \Delta\mathcal{A}_i} u_i(\alpha)$ for each i. □

The following lemma establishes some basic properties of B, C, and D. Elsewhere we take these properties for granted without reference to this lemma.

Lemma 11. *(i) W nonempty $\implies$ $D(W)$ nonempty $\implies$ $C(W)$ nonempty. (ii) W compact $\implies$ $D(W)$ compact $\implies$ $C(W)$ closed. (iii) $P(W)$ a nonempty line segment of slope -1 and $\underline{W}$ nonempty with $P(W \cup \underline{W}) = P(W)$ $\implies$ $B(W, \underline{W})$ nonempty; (iv) $\underline{W}$ also compact $\implies$ $B(W, \underline{W})$ compact $\implies$ co $B(W, \underline{W})$ compact.*

Proof. These results follow from the definitions of B, C, and D; the fact that the Nash equilibrium correspondence is upper hemicontinuous; and the fact that every stage game Nash equilibrium is enforced by a constant continuation value function. □

Proof of Lemma 2. Since W is a closed BSG set, and z^2 is the endpoint that most favors player 1, z_1^2 is the maximum payoff for player 1 that can be supported utilizing continuation values from W associated with the next period; i.e., $z_1^2 = \max\{v_1 \mid v \in B\big(C(W), D(W)\big)\}$. Because elements in $B(C(W), D(W))$ correspond to various disagreement points in $D(W)$, this maximization problem can be expressed using the bargaining solution (Eq. 7 with $C(W)$ in place of W), taking the disagreement point as the choice variable. Since $\text{level}(W) = \max_{w \in C(W)}(w_1 + w_2)$ for any given disagreement point $\underline{w} \in D(W)$, in bargaining player 1 earns the value $\underline{w}_1 + \pi_1(\text{level}(W) - \underline{w}_1 - \underline{w}_2)$. Therefore

$$
\begin{aligned}
z_1^2 = &\max_{\underline{w}, g, \alpha} \left(\pi_2 \underline{w}_1 - \pi_1 \underline{w}_2 + \pi_1 \,\text{level}(W)\right) \\
&\text{s.t. } \underline{w} = (1-\delta)u(\alpha) + \delta g(\alpha) \text{ and } g : \mathcal{A} \to W \text{ enforces } \alpha.
\end{aligned}
\tag{31}
$$

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 W is level(W), 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 $\underline{w}$ and combining terms, we see that Eq. 31 is equivalent to:

$$z_1^2 = \max_{\eta,\alpha} \big((1-\delta)\big(\pi_2 u_1(\alpha) - \pi_1 u_2(\alpha)\big) + \delta\big(\pi_2 z_1^2 - \pi_1 z_2^2\big) + \delta\eta(\alpha) + \pi_1 \operatorname{level}(W)\big),$$
$$\text{s.t.} \begin{cases} \eta : A \to [z_1^1 - z_1^2, 0] \text{ extended to } \Delta^{\mathrm{U}}A, \\ \alpha \in \Delta^{\mathrm{U}}A \text{ is a Nash equilibrium of } \big\langle A, (1-\delta)u + \delta(\eta, -\eta) + \delta z^2 \big\rangle. \end{cases} \tag{32}$$

In the objective function on the right, we substitute for z_2^2 using the equality $z_1^2 + z_2^2 = \operatorname{level}(W)$. Combining terms and moving the constant δz_1^2 to the left side yields

$$z_1^2 = \pi_1 \operatorname{level}(W) + \max_{\eta,\alpha} \left(\pi_2 u_1(\alpha) - \pi_1 u_2(\alpha) + \frac{\delta}{1-\delta}\eta(\alpha)\right), \tag{33}$$

subject to the conditions above. Recalling the definition of γ^2, we have the conclusion of the lemma. Analyzing endpoint z^1 in the same way gives the similar expression for z_2^1.

The optimum defining γ^i exists for each i because the stage game is finite, the set of feasible η functions is compact, and the Nash correspondence is upper hemicontinuous. Regarding the final claim of the lemma, let α^{NE} be a Nash equilibrium of the stage game. Note that α^{NE} and $\eta \equiv 0$ satisfy the conditions of the maximization problems that define γ^1 and γ^2, so $\gamma^2(d) \geq \pi_2 u_1(\alpha^{\mathrm{NE}}) - \pi_1 u_2(\alpha^{\mathrm{NE}})$ and $\gamma^1(d) \geq \pi_1 u_2(\alpha^{\mathrm{NE}}) - \pi_2 u_1(\alpha^{\mathrm{NE}})$. Summing, we have that $\gamma^2(d) + \gamma^1(d) \geq 0$. □

Proof of Lemma 3. Let z^1 and z^2 denote the endpoints of W and let Let z'^1 and z'^2 denote the endpoints of W'. Suppose that $\operatorname{span}(W) \geq \operatorname{span}(W')$. The larger span of W can support weakly more mixed actions in the stage game as equilibria than W' because it provides a greater range of continuation values. This comparison does not depend on the location of the endpoints or the levels of the two sets (which are mere constants in the players' payoffs), only their relative spans. Thus any mixed action that can be supported by W' can also be supported by W, so $\operatorname{level}(W) \geq \operatorname{level}(W')$. Furthermore, since each $\gamma^j(d)$ is increasing in d, Lemma 2 and the larger span of W imply that $z_i^j \geq z_i'^j$ for each $i \neq j$, which suffices to establish that W fully Pareto dominates W'. □

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.

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.

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

Geir B. Asheim. Individual and collective time-consistency. *Review of Economic Studies*, 64(3):427–443, July 1997.

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. Perfect equilibria in bargaining models. In Ken Binmore and Partha Dasgupta, editors, *The Economics of Bargaining*, chapter 5, pages 77–105. Basil Blackwell, Oxford, 1987.

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

Andreas Blume. Intraplay communication in repeated games. *Games and Economic Behavior*, 6(2): 181–211, March 1994.

Adam Brandenburger and Harborne Stuart. Biform games. *Management Science*, 53(4):537–549, April 2007.

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.

Sylvain Chassang. Building routines: Learning, cooperation, and the dynamics of incomplete relational contracts. *American Economic Review*, 100(1):448–465, March 2010.

Wouter J. den Haan, Garey Ramey, and Joel Watson. Job destruction and propagation of shocks. *The American Economic Review*, 90(3):482–498, 2000.

Wouter J. den Haan, Garey Ramey, and Joel Watson. Liquidity flows and fragility of business enterprises. *Journal of Monetary Economics*, 50:1215–1241, 2003.

Peter A. Diamond. Wage determination and efficiency in search equilibrium. *The Review of Economic Studies*, 49(2):217–227, April 1982.

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.

Katherine Doornik. Relational contracting in partnerships. *Journal of Economics and Management Strategy*, 15(2):517–548, Summer 2006.

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.

Yuk-fai Fong and Jin Li. Information revelation in relational contracts. Working paper, July 2010a. URL http://goo.gl/BfmMm.

Yuk-fai Fong and Jin Li. Relational contracts, efficiency wages, and employment dynamics. Working paper, April 2010b. URL http://goo.gl/Sqxn1.

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.

William Fuchs. Contracting with repeated moral hazard and private evaluations. *American Economic Review*, 97(4):1432–1448, September 2007.

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.

Garance Genicot and Debraj Ray. Bargaining power and enforcement in credit markets. *Journal of Development Economics*, 79:398–412, 2006.

Susanne Goldluecke and Sebastian Kranz. Infinitely repeated games with public monitoring and monetary transfers. Working paper, July 2010. URL http://goo.gl/hiw3Z.

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.

Marina Halac. Relational contracts and the value of relationships. Working paper, November 2010. URL http://goo.gl/nKSKg.

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

Benjamin E. Hermalin, Larry Li, and Tony Naughton. The welfare consequences of legal system improvement. Working paper, September 2011. URL http://goo.gl/pf2bD.

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.

Matthew O. Jackson and Thomas R. Palfrey. Efficiency and voluntary implementation in markets with repeated pairwise bargaining. *Econometrica*, 66(6):1353–1388, November 1998.

Matthew O. Jackson and Thomas R. Palfrey. Voluntary implementation. *Journal of Economic Theory*, 98 (1):1–25, May 2001.

Mikhail Klimenko, Garey Ramey, and Joel Watson. Recurrent trade agreements and the value of external enforcement. *Journal of International Economics*, 74(2):475–499, March 2008.

Sebastian Kranz and Susanne Ohlendorf. Renegotiation-proof relational contracts with side payments. Working paper, November 2010. URL http://goo.gl/iGCA3.

Ola Kvaløy and Trond E. Olsen. Team incentives in relational employment contracts. *Journal of Labor Economics*, 24(1):139–169, 2006.

Jonathan Levin. Multilateral contracting and the employment relationship. *Quarterly Journal of Economics*, 117(3):1075–1103, August 2002.

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

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

W. Bentley MacLeod and James M. Malcomson. Motivation and markets. *American Economic Review*, 88

(3):388–411, June 1998.

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.

Giuseppe Moscarini. Job matching and the wage distribution. *Econometrica*, 73(2):481–516, March 2005.

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

David G. Pearce and Ennio Stacchetti. The interaction of implicit and explicit contracts in repeated agency. *Games and Economic Behavior*, 23:75–96, 1998.

Christopher A. Pissarides. Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *The American Economic Review*, 75(4):676–690, 1985.

Christopher A. Pissarides. Search, wage bargains and cycles. *The Review of Economic Studies*, 54(3): 473–483, July 1987.

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

Horst Raff and David Schmidt. Cumbersome coordination in repeated games. *International Journal of Game Theory*, 29(1):101–118, 2000.

Garey Ramey and Joel Watson. Bilateral trade and opportunism in a matching market. *Contributions to Theoretical Economics*, 1(1), 2001.

Garey Ramey and Joel Watson. Contractual intermediaries. *Journal of Law, Economics, and Organization*, 18(2):362–384, October 2002.

Gary Ramey and Joel Watson. Contractual fragility, job destruction, and business cycles. *Quarterly Journal of Economics*, 112(3):873–911, August 1997.

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

Luis Rayo. Relational incentives and moral hazard in teams. *Review of Economic Studies*, 74(3):937–963, July 2007.

Ariel Rubinstein. Strong perfect equilibrium in supergames. *International Journal of Game Theory*, 9:1–12, 1979.

Ariel Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, January 1982.

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

Avner Shaked and John Sutton. Involuntary unemployment as a perfect equilibrium in a bargaining model. *Econometrica*, 52(6):1351–1364, 1984.

Jonathan Thomas and Tim Worrall. Self-enforcing wage contracts. *Review of Economic Studies*, 55(4): 541–553, October 1988.

Joel Watson. Contract and game theory: Basic concepts for settings with finite horizons. Working paper, January 2006. URL http://goo.gl/NFtle.

B Supplemental Appendix

B.1 Efficiency and bargaining power

Lemma 12. *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 γ^i_δ (see Eq. 10) 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

$$(1-\delta)u + \delta(\eta, -\eta) = \frac{1-\delta}{1-\delta'} \left((1-\delta')u + \delta'(\eta', -\eta')\right). \tag{34}$$

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. 10 at α, η, and δ is equal to the value at α, η', and δ'. These facts imply that $\gamma^i_\delta(d) = \gamma^i_{\delta'}(d')$, and so $\Gamma^\delta(d) = \Gamma^{\delta'}(d')$. □

Proof of Theorem 6 Fix $\hat{\delta}$. If $\Gamma_{\hat{\delta}}(\infty) > 0$ then there exists a number $\hat{d} > 0$ such that $\Gamma_{\hat{\delta}}(\hat{d}) > 0$. For any δ', let $d' \equiv \frac{\hat{\delta}(1-\delta')}{\delta'(1-\hat{\delta})} \hat{d}$. From Lemma 12 we see that $\Gamma_{\delta'}(d') = \Gamma_{\hat{\delta}}(\hat{d})$. As δ' converges to one, d' converges to zero, which implies that the maximal fixed point of Γ_δ is bounded below by $\Gamma_{\hat{\delta}}(\hat{d})$ for sufficiently high δ. Finally, it is clear from the definition of λ^* 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 $\Gamma_\delta(\infty) = 0$ 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 7 Let $a^* \in \arg\max_a \sum_i u_i(a)$. Under these strategies, the continuation value in state i is $\hat{z}^i \equiv u(a^i) + \pi \sum_j \left(u_j(a^*) - u_j(a^i)\right)$. Note that $\hat{z}^i$ does not depend on δ. Since $(\pi_2, -\pi_1) \cdot \left(u(a^2) - u(a^1)\right) > 0$, the payoff span $\hat{z}^2_1 - \hat{z}^1_1$ is strictly positive and constant in δ. We must check that the sequential rationality constraints are satisfied. Under disagreement in state i, player i is playing a stage game best response, and anticipates remaining in state i regardless of her action. Under agreement in either state, player i anticipates a loss of $\frac{\delta}{1-\delta} \frac{1}{2} (\hat{z}^2_1 - \hat{z}^1_1)$ if she deviates. Similarly, under disagreement in state $-i$, player i anticipates a loss of $\frac{\delta}{1-\delta} (\hat{z}^2_1 - \hat{z}^1_1)$ if she deviates. Hence efficiency is attained and all actions are sequentially rational for δ sufficiently high. Note that $\hat{z}^1$ and $\hat{z}^2$ are not necessarily the endpoints of $V(S_{\text{CE}})$. ■

Proof of Theorem 8 It suffices to restrict attention to the stage game and find a^1 and a^2 as described in Theorem 7. Choose a^{NE} satisfying the suppositions. For any small $\varepsilon' > 0$ let $a^{-i}_i(\varepsilon') \equiv a^{\text{NE}}_i + \varepsilon'$ and $a^i_i(\varepsilon') \equiv BR_i(a^i_{-i}(\varepsilon'))$. Near a^{NE}, since u is twice differentiable and a^{NE} is an equilibrium, for a_{-i} sufficiently close to a^{NE}_{-i} it follows that $\left|u_i\left(a^i_i(\varepsilon'), a_{-i}\right) - u_i\left(a^{-i}_i(\varepsilon'), a_{-i}\right)\right|$ is on the order of at most ε'^2. Since $dBR_i / da_{-i} < 1$, it follows that $a^{-i}_i - a^i_i > 0$ is on the order of at least ε'. Since $du_i / da_{-i} < 0$, for

a_i sufficiently close to a_i^{NE} it also follows that $u_i\big(a_i, a_{-i}^{-i}(\varepsilon')\big) - u_i\big(a_i, a_{-i}^{i}(\varepsilon')\big) > 0$ is on the order of at least ε'. Hence, for $\varepsilon' > 0$ sufficiently small, each player i strictly prefers a^{-i} to a^i. Since player i is best responding at a^i, the conditions of Theorem 7 are satisfied. ■

Proof of Theorem 9 Define

$$\hat{\Gamma}(d, \alpha^1, \alpha^2) \equiv \max_{\eta^1, \eta^2} \Big(\pi_2 \left(u_1(\alpha^1) - u_1(\alpha^2) \right) - \pi_1 \left(u_2(\alpha^1) - u_2(\alpha^2) \right) + \frac{\delta}{1-\delta} \left(\eta^1(\alpha^1) - \eta^2(\alpha^2) \right) \Big),$$
$$\text{s.t.} \begin{cases} \eta^1 : A \to \left[-d, 0\right], \text{ extended to } \Delta^{\text{U}} A, \\ \eta^2 : A \to \left[0, d\right], \text{ extended to } \Delta^{\text{U}} A, \\ \alpha^1 \in \Delta^{\text{U}} A \text{ is a Nash equilibrium of } \left\langle A, (1-\delta)u + \delta(\eta^1, -\eta^1) \right\rangle, \\ \alpha^2 \in \Delta^{\text{U}} A \text{ is a Nash equilibrium of } \left\langle A, (1-\delta)u + \delta(\eta^2, -\eta^2) \right\rangle. \end{cases} \tag{35}$$

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 $\Gamma(d)$, as is $\max_d \{d : d = \Gamma(d)\}$. ■

B.2 Generalization

This section proves Theorem 10 and Theorem 11.

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 $\hat{B}(X) \equiv \operatorname{co} B(C(X), D(X))$, $L(X) \equiv \max_{x \in X} \sum_{i=1}^n x_i$, and $\tilde{B}(X) \equiv \hat{B}(X) - \pi L(\hat{B}(X))$. The function $\tilde{B}$ normalizes the output of $\hat{B}$ 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.

Lemma 13. *Function $\hat{B}$ maps $\mathcal{X}$ to $\mathcal{X}$, and function $\tilde{B}$ maps $\mathcal{X}$ to $\mathcal{X}_0$. For every $\tilde{v} \in \hat{B}(X)$, it is the case that $\sum_{j=1}^n \tilde{v}_j = L(\hat{B}(X))$. Furthermore, for any $X \in \mathcal{X}$ and $x \in \mathbb{R}^n$, $\hat{B}(X + x) = \hat{B}(X) + \vec{\delta} * x$ 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 bargaining outcome is clearly continuous in the disagreement point and maximal joint value. Furthermore, $\tilde{B}(X)$ is a linear transformation of $\hat{B}(X)$ that makes every point balanced (joint value of zero). These facts imply the first part of the lemma. 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. 18, this is equivalent to replacing $\tilde{g}(\theta)$ with $\tilde{g}(\theta) + x$. Thus, the set of values that can be supported from the current period uniformly shifts by $\vec{\delta} * x$. Consequently, the set of bargaining outcomes likewise shifts. Finally, the level clearly changes as indicated. □

Hereinafter, all sets that we consider are understood to be compact subsets of $\mathbb{R}^n$. To establish the existence of a dominant BSG set, we first characterize the BSG sets—the fixed points of $\hat{B}$. To identify and compare BSG sets, we shall work with the function $\tilde{B}$. We start by demonstrating a relation between the fixed points of $\hat{B}$ and $\tilde{B}$. For each player i define $\phi_i \equiv \pi_i/(1-\delta_i)$, write $\phi = (\phi_1, \phi_2, \ldots, \phi_n)$, and let $\Phi \equiv \sum_{j=1}^{n} \phi_j$. Lemma 14 provides a linear relationship between the fixed points of $\hat{B}$ and $\tilde{B}$, in the direction ϕ.[22]

Lemma 14. *If $\tilde{W} \in \mathcal{X}$ and $X = \tilde{W} - \frac{\phi}{\Phi}L(\tilde{W})$, then $\tilde{W} = \hat{B}(\tilde{W})$ implies $X = \tilde{B}(X)$. If $X \in \mathcal{X}_0$ and $\tilde{W} = X + \phi L(\hat{B}(X))$, then $X = \tilde{B}(X)$ implies $\tilde{W} = \hat{B}(\tilde{W})$.*

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

$$\begin{aligned} \tilde{B}(X) &= \hat{B}(X) - \pi L(\hat{B}(X)) \\ &= \tilde{W} - \vec{\delta} * \frac{\phi}{\Phi} L(\tilde{W}) - \pi L(\tilde{W} - \vec{\delta} * \frac{\phi}{\Phi} L(\tilde{W})) \\ &= \tilde{W} - \vec{\delta} * \frac{\phi}{\Phi} L(\tilde{W}) - \pi L(\tilde{W})\Big(1 - \vec{\delta} \cdot \frac{\phi}{\Phi}\Big) = \tilde{W} - \Big(\pi + (\vec{1} - \pi)\vec{\delta} * \frac{\phi}{\Phi}\Big) L(\tilde{W}). \end{aligned} \tag{36}$$

The second line uses the property of $\hat{B}$ from Lemma 13 and that $V = \hat{B}(V)$. The third line uses the property of L from Lemma 13. Note that

$$\Phi - \vec{\delta} \cdot \phi = \sum_{j=1}^{n} \phi_j (1 - \delta_j) = \sum_{j=1}^{n} \frac{\pi_j}{1 - \delta_j}(1 - \delta_j) = 1. \tag{37}$$

Thus,

$$\tilde{B}(X) = \tilde{W} - \vec{\delta} * \frac{\phi}{\Phi} L(\tilde{W}) - \pi \frac{1}{\Phi} L(\tilde{W}) = \tilde{W} - (\vec{\delta} * \phi + \pi) \frac{1}{\Phi} L(\tilde{W}). \tag{38}$$

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

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

$$\begin{aligned} \hat{B}(\tilde{W}) &= \hat{B}\big(X + \phi L(\hat{B}(X))\big) \\ &= \hat{B}(X) + \vec{\delta} * \phi L(\hat{B}(X)) = \hat{B}(X) - \pi L(\hat{B}(X)) + \pi L(\hat{B}(X)) + \vec{\delta} * \phi L(\hat{B}(X)) \\ &= \tilde{B}(X) + (\pi + \vec{\delta} * \phi) L(\hat{B}(X)) = X + \phi L(\hat{B}(X)) = \tilde{W}. \end{aligned} \tag{39}$$

The second line uses the property of $\hat{B}$ from Lemma 13 and the third line uses the definition of $\tilde{B}$. The fourth line uses the assumption that $\tilde{B}(X) = X$ and that $\vec{\delta} * \phi + \pi = \phi$, which we showed above. □

We next show by construction that $\tilde{B}$ has a dominant fixed point. We start by identifying some properties of $\tilde{B}$.

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

Lemma 15. *The function $\tilde{B}$ is monotone: for every $X, X' \in \mathcal{X}$, $X \subset X'$ implies $\tilde{B}(X) \subset \tilde{B}(X')$. Furthermore, $\tilde{B}$ 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 $\tilde{B}(X^k)$ converges to $\tilde{B}(X)$.*

Proof. To prove the first part of the lemma, 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 $\tilde{w} \in B(C(X), D(X))$ there is an element $\underline{\tilde{w}} \in D(X)$ such that $\tilde{w} = \underline{\tilde{w}} + \pi\big(L(D(X)) - \vec{1} \cdot \underline{\tilde{w}}\big)$. Using the same disagreement point and choosing $\tilde{w}' = \underline{\tilde{w}} + \pi\big(L(D(X')) - \vec{1} \cdot \underline{\tilde{w}}\big)$, we have that $\tilde{w}' \in B(C(X'), D(X'))$. Thus, $\tilde{w} \in \hat{B}(X)$, $\tilde{w}' \in \hat{B}(X')$, and these points lie on the same line in direction π. Recall that $\tilde{B}$ normalizes by subtracting a vector proportional to π, so that resulting points have joint value of zero. This means that $\tilde{w} - \pi L(\hat{B}(X)) = \tilde{w}' - \pi L(\hat{B}(X'))$, which proves $\tilde{B}(X) \subset \tilde{B}(X')$.

To prove the second part of the lemma, 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 Lemma 13. Thus, $\hat{B}(X^k)$ converges to $\hat{B}(X)$ and so $\tilde{B}(X^k)$ converges to $\tilde{B}(X)$. □

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

Lemma 16. *For $X, X' \in \mathcal{X}_0$, $X \subset X'$ implies that $L(\hat{B}(X)) \leq L(\hat{B}(X'))$.*

Proof of Theorem 10 To construct a dominant fixed point of $\tilde{B}$, start with a large element of $\mathcal{X}_0$ that is guaranteed to be a superset of any fixed point. Since Θ and $\mathcal{A}$ are finite, there exists $K \in \mathbb{R}_+$ such that all stage-game payoffs are bounded below by $-K(1 - \max_j \delta_j)$ and above by $K(1 - \max_j \delta_j)$. Let $X^1 \equiv \{x \in \mathbb{R}_0^n \mid -K \leq x \leq K \text{ for all } i\}$. Then every fixed point of $\tilde{B}$ is a subset of X^1 and that $\tilde{B}(X^1) \subset X^1$. Define the sequence $\{X^k\}$ inductively by $X^{k+1} \equiv \tilde{B}(X^k)$, for all $k > 1$. Since $\tilde{B}$ 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.

Lemma 15 implies that $X^* = \tilde{B}(X^*)$. To see this, note that $X^* \subset X^{k+1} = \tilde{B}(X^k)$ for all k. Because $\tilde{B}$ is continuous on decreasing sequences, $\tilde{B}(X^k)$ converges to $\tilde{B}(X^*)$ and so we have that $X^* \subset \tilde{B}(X^*)$. In addition, because $\tilde{B}$ is monotone and $X^* \subset X^k$, we have $\tilde{B}(X^*) \subset \tilde{B}(X^k) = X^{k+1}$ for all k. That X^{k+1} converges to X^* then implies that $\tilde{B}(X^*) \subset X^*$.

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

Thus, we have established that X^* is a fixed point of $\tilde{B}$ and every other fixed point of $\tilde{B}$ is contained in X^*. Define $\tilde{W}^* \equiv X^* + \phi L(\hat{B}(X^*))$. We finish the proof by showing that $\tilde{W}^*$ is a BSG set for the game and it dominates every other BSG set. That $\tilde{W}^*$ is an BSG set follows immediately from Lemma 14. For the second step, consider any other BSG set $\tilde{W}$ and define $X \equiv \tilde{W} - \frac{\phi}{\Phi} L(\tilde{W})$. From Lemma 14 we know that X is a fixed point of $\tilde{B}$. We also know that $X \subset X^*$.

From the relationship between fixed points of $\underline{w}^1$ and $\tilde{B}$, we know that $\tilde{W} = X + \phi L(\underline{w}^1(X))$. Take any $\tilde{v} \in \tilde{W}$ and let $\tilde{x} \in X$ be such that $\tilde{v} = \tilde{x} + rL(\underline{w}^1(X))$. Since $X \subset X^*$, we have that $\tilde{x} \in X^*$ and thus $\tilde{v}' \equiv \tilde{x} + \phi L(\underline{w}^1(X^*)) \in \tilde{W}$. Comparing $\tilde{v}$ and $\tilde{v}'$, we see that $\tilde{v}' - \tilde{v} = \phi\big(L(\underline{w}^1(X^*)) - L(\underline{w}^1(X))\big)$. From Lemma 16, we know that $L(\underline{w}^1(X^*)) \geq L(\underline{w}^1(X))$. 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 $\tilde{W}^*$ dominates $\tilde{W}$. ■

Proof of Theorem 11 We first work through the construction of $\tilde{z}^2$, which is player 1's most preferred point in $\tilde{W}^*$. Everything is analogous for $\tilde{z}^1$. In this environment, Eq. 31 becomes:

$$\begin{aligned} \tilde{z}_1^2 = \max_{\underline{\tilde{w}}, \tilde{g}, \alpha} & \left(\pi_2 \underline{\tilde{w}}_1 - \pi_1 \underline{\tilde{w}}_2 + \pi_1 \operatorname{level}(\tilde{W}^*)\right), \\ \text{s.t.} & \begin{cases} \underline{\tilde{w}} = \sum_{\theta \in \Theta} f(\theta|\alpha)\big(u(\alpha, \theta) + \vec{\delta} * \tilde{g}(\theta)\big), \\ \tilde{g} : \Theta \to \tilde{W}^* \text{ enforces } \alpha. \end{cases} \end{aligned} \tag{40}$$

Define $\eta(\theta) \equiv \tilde{g}_1(\theta) - \tilde{z}_1^2$ and $\hat{u}_i(\alpha) \equiv \sum_{\theta \in \Theta} f(\theta|\alpha) u_i(\alpha_i, \theta)$. After rearranging terms as before, the optimization problem of Eq. 32 becomes:

$$\begin{aligned} \tilde{z}_1^2 = \max_{\eta, \alpha} & \left(\big(\pi_2 \hat{u}_1(\alpha) - \pi_1 \hat{u}_2(\alpha)\big) + \big(\pi_2 \delta_1 \tilde{z}_1^2 - \pi_1 \delta_2 \tilde{z}_2^2\big) + \big(\pi_2 \delta_1 + \pi_1 \delta_2\big) \hat{\eta}(\alpha) + \pi_1 \operatorname{level}(\tilde{W}^*)\right), \\ \text{s.t.} & \begin{cases} \eta : \Theta \to \left[\tilde{z}_1^1 - \tilde{z}_1^2, 0\right], \text{ with } \hat{\eta}(\alpha) \equiv \sum_{\theta \in \Theta} f(\theta|\alpha) \eta(\theta), \\ \alpha \in \Delta^{\mathrm{U}} \mathcal{A} \text{ is a Nash equilibrium of } \left\langle \mathcal{A}, \hat{u} + \vec{\delta} * \big((\hat{\eta}, -\hat{\eta}) + \tilde{z}^2\big) \right\rangle. \end{cases} \end{aligned} \tag{41}$$

Substituting in ψ, using $\tilde{z}_1^2 + \tilde{z}_2^2 = \operatorname{level}(\tilde{W}^*)$, and rearranging terms yields:

$$\tilde{z}_1^2 = \frac{\pi_1(1 - \delta_2)}{1 - \psi} \operatorname{level}(\tilde{W}^*) + \tilde{\gamma}^2(\operatorname{span}(\tilde{W}^*)), \tag{42}$$

where γ^2 is defined in Theorem 11. Similar calculations yield:

$$\tilde{z}_2^1 = \frac{\pi_2(1 - \delta_1)}{1 - \psi} \operatorname{level}(\tilde{W}^*) + \tilde{\gamma}^1(\operatorname{span}(\tilde{W}^*)). \tag{43}$$

Summing these expressions, we have:

$$\tilde{z}_2^1 + \tilde{z}_1^2 = \operatorname{level}(\tilde{W}^*) + \tilde{\gamma}^1(\operatorname{span}(\tilde{W}^*)) + \tilde{\gamma}^2(\operatorname{span}(\tilde{W}^*)). \tag{44}$$

Using $\tilde{z}_2^1 + \tilde{z}_1^1 = \operatorname{level}(\tilde{W}^*)$ to substitute for $\tilde{z}_2^1$, we obtain $\operatorname{span}(\tilde{W}^*) = \tilde{\gamma}^1(\operatorname{span}(\tilde{W}^*)) + \tilde{\gamma}^2(\operatorname{span}(\tilde{W}^*))$, so we conclude that $\operatorname{span}(\tilde{W}^*)$ is a fixed point of $\tilde{\Gamma}$ as in the basic model.

The construction of $\tilde{W}^*$ proceeds as in the basic model. We first find the maximal fixed point of $\tilde{\Gamma}$. We then have to calculate level$(\tilde{W}^*)$, 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:

$$\text{level}(\tilde{W}^*) = \max_{\tilde{g},\alpha}\Big(\hat{u}_1(\alpha) + \hat{u}_2(\alpha) + \sum_{\theta\in\Theta} f(\theta|\alpha)\vec{\delta}\cdot\tilde{g}(\theta)\Big) \text{ s.t. } \tilde{g}:\Theta\to\tilde{W}^* \text{ enforces } \alpha, \tag{45}$$

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 $\tilde{W}^*$.

With the same steps taken above, we rewrite this maximization problem by substituting for $\tilde{g}$ using the function η, where we have $\tilde{g}_1(\theta) = \eta(\theta) + \tilde{z}_1^2$ and $\tilde{g}_2(\theta) = \tilde{z}_2^1 - \text{span}(\tilde{W}^*) - \eta(\theta)$. This yields:

$$\text{level}(\tilde{W}^*) = \delta_1\tilde{z}_1^2 + \delta_2\tilde{z}_2^1 - \delta_2\,\text{span}(\tilde{W}^*) + \chi(\text{span}(\tilde{W}^*)), \tag{46}$$

where

$$\begin{aligned}\chi(\tilde{d}) = \max_{\eta,\alpha}\;& \hat{u}_1(\alpha) + \hat{u}_2(\alpha) + (\delta_1-\delta_2)\hat{\eta}(\alpha)\\ \text{s.t.}\;& \begin{cases}\eta:\Theta\to[-\tilde{d},0], \text{ with } \hat{\eta}(\alpha)\equiv\sum_{\theta\in\Theta} f(\theta|\alpha)\eta(\theta),\\ \alpha\in\Delta^{\mathrm{U}}\mathcal{A} \text{ is a Nash equilibrium of } \big\langle\mathcal{A}, \hat{u}+\vec{\delta}*(\hat{\eta},-\hat{\eta})\big\rangle.\end{cases}\end{aligned} \tag{47}$$

We can use Eq. 42 and Eq. 43 to substitute for $\tilde{z}_1^2$ and $\tilde{z}_2^1$ in Eq. 46. After solving for level$(\tilde{W}^*)$ and simplifying, we obtain:

$$\text{level}(\tilde{W}^*) = \frac{1-\psi}{(1-\delta_1)(1-\delta_2)}\begin{pmatrix}\delta_1\tilde{\gamma}^2(\text{span}(\tilde{W}^*)) + \delta_2\tilde{\gamma}^1(\text{span}(\tilde{W}^*))\\ -\;\delta_2\,\text{span}(\tilde{W}^*) + \chi(\text{span}(\tilde{W}^*))\end{pmatrix}. \tag{48}$$

The foregoing argument proves the first and second parts of the theorem. The third part follows from Eq. 42 and Eq. 43, and that the joint value is level$(\tilde{W}^*)$. ■