An Active-Contracting Perspective on Equilibrium Selection in Relational Contracts

David A. Miller and Joel Watson*

January 8, 2023

Abstract

In Miller and Watson (2013) we introduced *contractual equilibrium* for repeated games with bargaining. In this article, using a simple principal-agent model as an example, we highlight how contractual equilibrium resolves conceptual problems with the prior literature's standard approach to equilibrium selection in relational contracting games. For instance, typical models give the principal the right to make a take-it-or-leave-it offer in each period, yet in the equilibria typically selected the principal does not have bargaining power to exert. More broadly, we show that this and other modeling features that seem to impose substantive structure actually do not affect the set of equilibrium payoffs in the model, absent an equilibrium selection by the analyst. Contractual equilibrium, in contrast, axiomatizes how the parties should interpret their bargaining outcomes, putting the equilibrium selection in their hands. In contractual equilibrium, the agent's effort is increasing in her own bargaining power, and the way the parties play under disagreement depends on the history of the agent's actions.

Keywords: Equilibrium selection, active contracting, bargaining power, relationships
JEL codes: C73, C78, D23, D86

*Miller: University of Michigan, Ann Arbor, USA, econdm@umich.edu; Watson (corresponding author): University of California, San Diego, USA, jwatson@ucsd.edu. The authors thank the editors and anonymous reviewers for their generous comments and encouragement.

1 Introduction

A relational contract is a plan for how the parties in a productive relationship should behave, such that all parties agree to the plan and none has an incentive to deviate from it, even though not every aspect of the plan can be enforced by an external legal authority. In the relational-contracting literature, productive relationships are modeled as repeated games, where the requirement that no party has an incentive to deviate implies that the plan must form a subgame perfect equilibrium (or a refinement thereof). However, repeated games tend to have many equilibria, naturally raising the question of which equilibrium the parties should be expected to select. This question distinguishes the relational-contracting literature from the larger repeated-games literature: while the repeated-games literature is largely concerned with characterizing the set of equilibria, the relational-contracting literature focuses on equilibria that jointly benefit the contracting parties.

In this paper we address conceptual problems with what has been the standard approach to modeling and equilibrium selection in the relational-contracting literature. Most importantly, the standard approach provides no way to represent how the parties meaningfully discuss what relational contract they should agree on, and there is no role for parties to exert bargaining power when negotiating over a contract. An implication is that common modeling elements, such as ultimatum offers, externally enforced wages, and automatic separation on disagreement, are effectively irrelevant, exerting influence over neither the set of equilibrium payoffs nor the selection from that set.

In 2013 we published "A Theory of Disagreement in Repeated Games with Bargaining" (Miller and Watson 2013), proposing the solution concept of *contractual equilibrium* for relational contracting. Our motivation was to directly address how players can discuss and bargain over how to play in repeated games. Contractual equilibrium is an axiomatic approach to equilibrium selection based on the idea that statements players make in the negotiation process serve to coordinate their future behavior. Bargaining entails meaningful discussion, and exercise of bargaining power arises as an implication. The players can bargain not only at the start of the relationship, but also repeatedly as the relationship unfolds. In this article we explain how the contractual-equilibrium approach can simplify and regularize the study of relational contracts, bringing new insights and stronger interpretation. We argue that contractual equilibrium can expand the scope of relational-contracting analysis. Also, on the technical side, we describe how our theory extends to settings with outside options.

We illustrate our approach starting with a widely known model of a principal and an agent, a perfect-monitoring special case of Levin (2003). This model features bargaining between parties in the form of a take-it-or-leave-it offer of a wage and bonus, which would seem to allocate all bargaining power to the party that makes the offer. However, reallocating the power to make the offer or even eliminating the bargaining phase does not change the set of subgame perfect equilibrium payoffs. Levin's model also features both externally enforced and voluntary transfers at the end of the period, but we show that the equilibrium payoff set is unchanged by condensing all transfers into a voluntary phase immediately after the bargaining phase. Finally, Levin's model assumes that if the agent rejects the principal's offer then mechanically they must get their outside options for the period, but we show that the equilibrium payoff set is unchanged if the outside options can be triggered voluntarily by either party. These observations imply that giving one party (e.g., the principal) the power to make offers, restricting the parties to take their outside options under disagreement, and making the wage externally enforceable play no substantive roles in the analysis.

With regard to equilibrium selection, Levin (2003) and most of the subsequent literature have largely adopted an intuitive approach, selecting equilibria on the Pareto frontier of the equilibrium payoff set, because "it is natural to focus on contracts that maximize the parties' joint surplus" (Levin 2003, p. 840). In our view, this intuitive approach takes the selection decision out of the hands of the contracting parties, leaving no role for meaningful discussion or exercise of bargaining power. The bargaining phase in particular serves no real purpose; it merely gives the players an opportunity to state their intention to play the equilibrium that the analyst has already selected. Moreover, this opportunity does not imply that they must state their intentions truthfully: because the proposer's offer is cheap talk,[1] the parties can simply act according to the selected equilibrium regardless of what is said during bargaining.

1.1 The axioms of contractual equilibrium

To model how the parties jointly decide on their relational contract—that is, to model *active contracting*—an axiomatic approach is most appropriate. In our view, an axiom that restricts play to the Pareto frontier directly is tantamount to assuming the conclusion. Instead, axioms should generate endogenously meaningful bargaining, connecting what the parties choose to say in the bargaining phase to how they behave in the continuation game. If the parties reach an agreement that is credible, they should follow through on it. Moreover, in Miller and Watson (2013) and Watson (2013) we show that substantive refinement also requires a "theory of disagreement" that constrains behavior after a responder rejects

[1] Recall that the assumption that the wage is externally enforced at the end of the period is irrelevant because it may as well be paid immediately.

a proposer's offer. In Levin (2003), the theory of disagreement is that the parties receive their outside options for the period. In contrast, contractual equilibrium employs a more nuanced theory of disagreement, which allows the players to remain in the relationship. If they disagree in the current period, contractual equilibrium specifies that they fall back on behavior that they jointly arranged in their most recent agreement,[2] and this behavior may be responsive to their history.

To state the axioms simply, we first simplify the model to have, in each period, the following structure: One of the parties is randomly selected to be the proposer and makes a take-it-or-leave-it offer, the responder can agree or disagree, voluntary transfers are exchanged, the parties simultaneously decide whether to trigger the outside option for one period, and the agent chooses effort to exert (if the outside option was not triggered). The proposer's offer can be taken as a suggested continuation payoff vector, which incorporates an immediate wage or bonus along with expected future effort choices and payments. This model has the same set of subgame perfect equilibrium payoffs as in Levin (2003).

Contractual equilibrium imposes three axioms on how offers and responses in the bargaining stage are interpreted. First, the *Internal Agreement Consistency* (IAC) axiom specifies that if the proposer offers a continuation payoff that is available at an alternative history in the same equilibrium, and the responder agrees, then this is considered a credible agreement and they should switch to playing as if at that alternative history. Second, the *No-Fault Disagreement* (NFD) axiom specifies that if the responder rejects the proposer's offer, then their continuation play should depend neither on which player was selected to be the proposer, nor on what proposal was offered, nor on what voluntary transfer was made.[3] Together, IAC and NFD endow the bargaining phase with a well-defined disagreement point and bargaining set, enabling a standard bargaining analysis based on the Nash bargaining solution: The parties choose the maximal joint payoff from the bargaining set, through a transfer the proposer gets all the surplus relative to the disagreement point, and from the start of the period the expected outcome splits the surplus in proportion to the likelihood of being selected to propose.[4]

[2]If they have never reached an agreement, then the fallback behavior must be specified from outside the relationship, such as by a social norm.

[3]In writing Miller and Watson (2013) we struggled to determine the appropriate enforcement environment regarding the transfer stage, and considered two possibilities: either the agreement involves an externally enforced contract (such as a signed bank check) that causes the transfer to be made; or that while an agreement may be reached in the negotiation phase, the "deal is sealed" only once the agreed-upon transfer is made in the transfer phase. We later realized that all we need is the simpler idea that play under disagreement should be invariant to any voluntary transfer that may have been made.

[4]Miller and Watson (2013) shows that the bargaining phase can be generalized to allow for arbitrarily many rounds of back-and-forth bargaining, with random proposer selection and random breakdown.

The third axiom is *Pareto External Agreement Consistency* (PEAC). It specifies that if the proposer offers a continuation payoff that is on the Pareto frontier among equilibria that satisfy IAC and NFD, and the responder agrees, then this is considered a credible agreement and they should switch to playing as if in that alternative equilibrium. PEAC allows the contracting parties to select among the best equilibria that satisfy IAC and NFD.

Unlike the intuitive approach of simply selecting the Pareto frontier, contractual equilibrium puts the decision in the hands of the contracting parties, by endowing their agreements and disagreements with meaning. Under the axioms, prepending a bargaining phase onto the stage game is no longer irrelevant—with meaningful discussion, it gives the parties the opportunity to plan out their relationship, thereby eliminating equilibria that are inconsistent with their joint incentives and relative bargaining weights (defined as the probabilities of being selected to propose).

1.2 Insights from contractual equilibrium

In addition to conceptual clarity, contractual equilibrium enables a straightforward analysis and generates new insights. In Miller and Watson (2013) we showed that contractual equilibrium identifies a unique joint payoff, which may be less than the highest feasible in the game, even if the parties are very patient. The theory thus provides a critical insight: We can predict that the parties will achieve a continuation value on a relevant Pareto frontier *only if* we model them as actively contracting, and this necessarily involves exercise of bargaining power, with attendant effects on the attainable joint value. What the parties can attain depends on their relative bargaining power and the properties of the stage game.[5] We showed how to construct a contractual equilibrium in an arbitrary two-player repeated game, by solving three interrelated optimization problems.

Figure 1 illustrates the comparison between subgame perfect equilibrium and contractual equilibrium in the simple principal-agent relationship, while also demonstrating the practical implications of our theory. The curve extending from the origin gives the payoff vector of production at various effort levels, where the agent's negative payoff represents the disutility of effort and the principal's positive payoff is the benefit of effort to the firm. The origin represents zero effort. As the agent's effort level increases, the principal's payoff rises and the agent's payoff falls. Points shown in the positive quadrant are continuation values achievable with positive effort and transfers from the principal to the agent.

[5] Halac (2012) studies a different kind of principal-agent model in which bargaining power also matters due to private information, also explained in Halac (2022). In that model, the bargaining power creates asymmetric information rents that determine the extent to which the principal's private information about the value of the relationship can be signaled or screened.

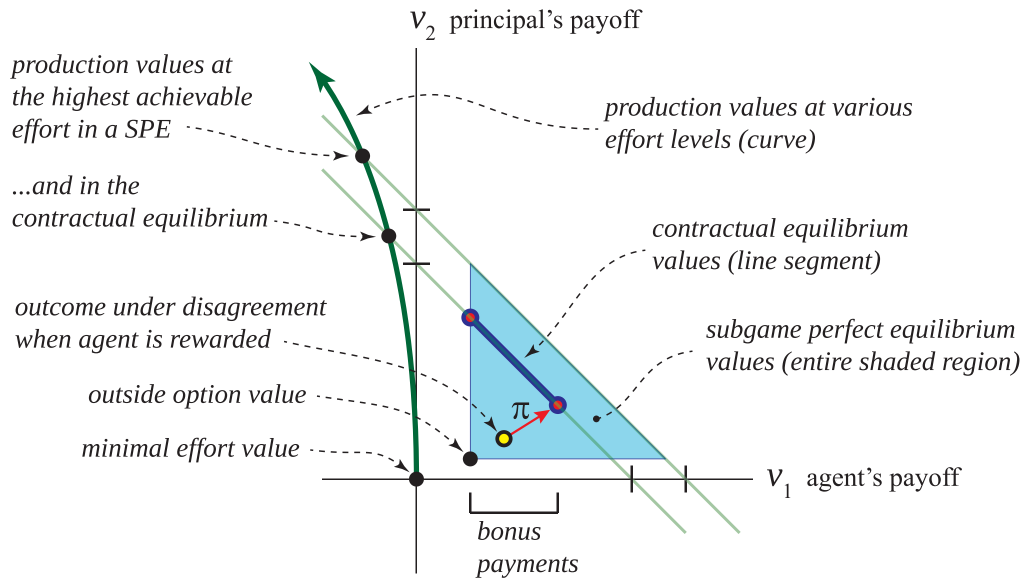

FIGURE 1. Illustration of the analysis.

As we show in this article, the standard analysis of the literature prior to Miller and Watson (2013) leads to the set of subgame perfect equilibrium values, which appears as the shaded triangular region in the figure, with no further refinement. The standard analysis supports equilibria with continuation values close to the outside-option value as well as equilibria with higher joint values. Further, for a given joint value above the outside option, there are equilibria that give much of the surplus to the agent and other equilibria that give much of the surplus to the principal. Exogenous bargaining power plays no role.

In contrast, the set of contractual-equilibrium values is the thick line segment indicated in Figure 1, within the shaded region. Active negotiation implies that every equilibrium continuation under agreement has the same joint value. Further, the principal's lowest equilibrium payoff is not the principal's outside-option value; rather, it is the value obtained when the players renegotiate from the disagreement point that most favors the agent. In the case shown in the figure, this disagreement point, labelled "outcome under disagreement when agent is rewarded," entails the players taking the outside option in the current period and coordinating on the principal's worst equilibrium value from the next period. Renegotiation leads to the highest attainable equilibrium joint value. The negotiation surplus is shared in proportion to the probabilities of being chosen to make the offer, π_1 for the agent and π_2 for the principal. Thus, the agreement value, disagreement point, and outside-option value are on a line segment in the direction $\pi = (\pi_1, \pi_2)$, the vector of bargaining weights.[6]

When the model's predictions are translated into contractual specifications in the real world, the agent's outside-option payoff corresponds to the wage, which must be paid even under disagreement to prevent the agent from quitting the relationship to take up employment elsewhere. Furthermore, the horizontal length of the contractual-equilibrium value set corresponds to the maximal discretionary bonus that is paid to reward the worker for exerting high effort (following histories when the agent should be rewarded). Essentially, the principal's bargaining power creates a hold-up problem that binds the discretionary bonuses. If the principal's bargaining weight π_2 were increased, the vector π would have a steeper slope in Figure 1, implying that the contractual-equilibrium line segment would shorten. The line segment would also shift closer to the origin, because the reduction of equilibrium bonuses lowers the worker's effort incentives.

1.3 Plan of the paper

This article proceeds as follows. In the next section we establish our negative results on the prior literature's standard approach to studying relational contracts. We do this by formally comparing the Levin (2003) model to several variants, including one that strips out much of its structure and has no negotiation element. In Section 3, we describe the axiomatic foundation of contractual equilibrium and explain how it works in practice. We also note how to reinterpret contractual equilibria in the repeated game as reduced-form equilibria in a hybrid game, which in each period features first a cooperative bargaining phase with transfers, and then a non-cooperative stage game. The hybrid game is simpler to analyze, and generates an equivalent set of equilibria. Section 4 provides more details of the equilibrium construction in the principal-agent application. In the Conclusion, we make additional comments on conceptual issues and connections with the literature, discuss extensions to imperfect monitoring (studied in Miller and Watson 2013) and imperfect external enforcement (studied in Watson, Miller, and Olsen 2020), and note open questions and opportunities for future work.

[6] There are some variations in the disagreement point when the agent is to be rewarded, depending on the relation between the outside-option value and the minimal-effort value. To simplify, for much of the analysis we assume that these are equal. The disagreement point when the agent is to be punished entails the agent exerting the highest effort sustainable in equilibrium, implying zero bargaining surplus.

2 Irrelevant features of familiar relational contracting models

The basic relational-contracting framework has an infinite time horizon, discrete periods, and a fixed production technology that the players engage in repeatedly. The time period is denoted $t = 1, 2, 3, \ldots$. Players discount future payoffs according to a common discount factor $\delta \in (0, 1)$.

To make our points most directly, we focus on the Levin (2003) model of relational contracting, one of the standard models in the literature. There are two parties: the agent whom we call player 1, and the principal whom we call player 2. We generalize Levin's model to allow for random-proposer bargaining, but we specialize to the environment without private information and with perfect monitoring, and with a specific quadratic cost function. Our conclusions do not depend on these specializations; in Miller and Watson (2013) we show how our approach generalizes to arbitrary repeated games with imperfect public monitoring, which also implicitly nests the kind of private information modeled by Levin.

Model 1 (Levin)**.** *In each period, the agent and principal interact in the following three phases, in the order shown:*

- **Negotiation phase:** *Nature selects one of the players to be the proposer. The agent is selected with fixed probability π_1 and the principal is selected with probability π_2, where $\pi_1, \pi_2 \geq 0$ and $\pi_1 + \pi_2 = 1$.*[7] *The proposer offers a spot contract to the responder, consisting of an externally enforced net monetary transfer $m = (m_1, m_2)$ to take place in the period and a suggested schedule of voluntary bonuses that the players may exchange at the end of the period. The transfer satisfies $m_1 + m_2 \leq 0$ and can be interpreted as a wage payment, along with "money burning" or transfers to a third party if $m_1 + m_2 < 0$. The responder either accepts or rejects the offer.*
- **Production phase (stage game):** *If the responder accepted in the negotiation phase, then the players are engaged for the period, the agent selects an effort level $e \geq 0$, and the players receive stage-game payoff vector $u = (-e^2, e + e^2)$. If the responder rejected in the negotiation phase, then the players are disengaged for the period, and they receive their outside-option payoffs given by vector $\underline{u}$.*
- **Bonus phase:** *If they are engaged, simultaneously the players make voluntary monetary transfers and can also burn money. Let $b = (b_1, b_2)$ denote the vector of total monetary transfers received, where $b_1 + b_2 \leq 0$.*[8]

[7] Levin assumes $\pi_2 = 1$.

[8] Operationally, each player chooses a non-negative amount to send to the other player, and also chooses a non-negative amount to burn. Since only the net transfers are important, for the rest of this article we ignore how simultaneous voluntary transfers are operationalized.

Total payoffs in the period are the sum of the stage-game payoffs and monetary transfers, normalized by $(1-\delta)$*, so that player* i *obtains* $(1-\delta)(m_i+u_i+b_i)$ *if engaged and* $(1-\delta)\underline{u}_i$ *if disengaged.*

With the particular form of production-phase payoffs assumed here, $-e^2$ for the agent and $e + e^2$ for the principal, the joint value of production conveniently equals the effort level e, while agent has a standard convex cost. Although effort is unbounded, due to the quadratic cost and time discounting it turns out that the players cannot sustain arbitrarily high effort. We also assume $\underline{u}_1 + \underline{u}_2 \geq 0$, so the joint value of receiving the outside option in a period exceeds the joint value of producing at minimal effort. The normalization allows us to interpret the present value of a series of payoffs as the per-period average.

This modeling approach incorporates several special features that at first glance seem to impose structure with substantive implications, but turn out to be irrelevant when analyzing the model using subgame perfect equilibrium:

1. The inclusion of the bonus phase, which gives the parties a second opportunity to make monetary transfers in the period;

2. The assumption that the monetary transfers at the beginning of the period are externally enforced if the offer in the negotiation phase is accepted;

3. The assumption that the outside option is inextricably triggered if the offer in the negotiation phase is rejected;

4. The modeling of the negotiation phase using an ultimatum offer;

5. The inclusion of a recommended bonus schedule in the ultimatum offer.

The voluntary bonuses combined with the externally enforced wages suggest an environment of incomplete external enforcement, but in fact no external enforcement is needed to generate the same conclusions. The outside option directly constrains the set of equilibrium values, but in a subgame perfection analysis it does not matter whether it is triggered automatically upon rejection or voluntarily. The ultimatum offer may seem to endow the proposer with bargaining power because it forces the responder to choose between the offered terms and the outside option, but in a subgame perfection analysis the privilege to make ultimatum offers is of no consequence.

In fact, the set of payoffs attainable in subgame perfect equilibrium in Model 1 is identical to that of the Model 2, below, in which there is no negotiation and no external enforcement, and the outside option is literally optional:

Model 2 (Basic)**.** *In each period, the agent and principal interact in the following two phases, in the order shown:*

- **Transfer phase:** *The players simultaneously pay voluntary monetary transfers and can also burn money. Let* $m = (m_1, m_2)$ *denote the vector of net monetary transfers received, where* $m_1 + m_2 \leq 0$.
- **Production phase (stage game):** *The players simultaneously decide whether to engage or disengage in the current period. If either player elects to disengage, then there is no production in the period and the players receive their outside-option payoffs given by vector* $\underline{u}$. *If they are engaged, the agent selects an effort level* $e \geq 0$ *and the players receive stage-game payoff vector* $u = (-e^2, e + e^2)$.

Total payoffs in the period are the sum of the stage game payoffs and monetary transfers, normalized by $(1-\delta)$, *so that player* i *obtains* $(1-\delta)(m_i+u_i)$ *if engaged and* $(1-\delta)(m_i+\underline{u}_i)$ *if disengaged.*

In both models, to bound the joint value that can be generated in equilibrium, we must find the highest effort level that the agent could be given the incentive to choose. It is useful to define $g(\sigma)$ as the highest value of e that satisfies either $(1-\delta)e^2 \leq \delta(e-\sigma)$ or $e = 0$, which is an incentive condition explained as follows. Suppose σ is the sum of the players' outside-option payoffs. If player 1 is supposed to choose effort $e > 0$ in the production phase of a given period, then the best deviation would be to choose zero effort, yielding a gain of $(1-\delta)e^2$ in the current period. If e were the highest joint value that the players could achieve in the continuation game from the next period, then the greatest punishment that could be imposed on player 1 is no more than $\delta(e-\sigma)$ in current-period terms. A little algebra reveals that

$$g(\sigma) = \frac{\delta}{2(1-\delta)}\left(1 + \sqrt{1 - 4\left(\frac{1-\delta}{\delta}\right)\sigma}\right) \qquad \text{for all } \sigma \leq \delta/4(1-\delta),$$

while $g(\sigma) = 0$ for all $\sigma > \delta/4(1-\delta)$. Note that g is a decreasing function and $g(0) = \delta/(1-\delta)$.

Proposition 1. *In both Model 1 and Model 2, under subgame perfect equilibrium (SPE):*

i. The highest effort level attainable is $e^* = g(\underline{u}_1 + \underline{u}_2)$;

ii. If $\underline{u}_1 + \underline{u}_2 \leq \delta/4(1-\delta)$ *then the set of equilibrium continuation values is*

$$V^{\text{SPE}} = \{v \in \mathbb{R}^2 \mid v_1 \geq \underline{u}_1, v_2 \geq \underline{u}_2, v_1 + v_2 \leq e^*\},$$

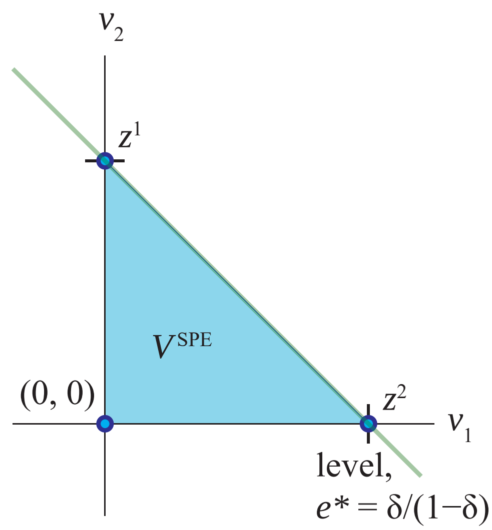

FIGURE 2. Principal-agent SPE set for the case of $\underline{u}_1 = \underline{u}_2 = 0$.

whereas if $\underline{u}_1 + \underline{u}_2 > \delta/4(1-\delta)$ *then it is* $V^{\text{SPE}} = \{\underline{u}\}$.

The set V^{SPE} is pictured in Figure 2 for the case of $\underline{u}_1 = \underline{u}_2 = 0$. We present this case for comparison with a set we construct in Section 4.

Proof. Consider first the case of $\underline{u}_1+\underline{u}_2 \leq \delta/4(1-\delta)$ and Model 2. It is easy to confirm that the following specification of strategies is a SPE that achieves continuation value $\underline{u}$. In every period, regardless of the history, the players are supposed to make no transfers. Regardless of the history and the transfers actually made, the players are supposed to disengage in the current period. In the event that production occurs (due to both players deviating to engage rather than disengage), then the agent is supposed to select $e = 0$.

We next describe a SPE that achieves continuation value $z^2 = (e^* - \underline{u}_2, \underline{u}_2)$. So long as there have been no deviations in the past, the strategy profile prescribes that in the transfer phase, the principal transfers $e^* + (e^*)^2 - \underline{u}_2$ to the agent whereas the agent transfers nothing. After the prescribed transfer, the players are supposed to remain engaged so that production occurs, and the agent is supposed to choose effort e^*. If there are any deviations, continuation play reverts to the strategy specification described above that achieves continuation value $\underline{u}$ from the beginning of the period; that is, there are no transfers, the players disengage, and the agent exerts zero effort. In every contingency, the continuation values of the prescribed behavior weakly improve on those of the punishment path, and therefore the strategy profile is a SPE.

All other values in $v \in V^{\mathrm{SPE}}$ can be achieved in equilibrium using appropriately prescribed first-period transfers and then conditioning the continuation as described to achieve z^2 from the next period. In particular, continuation value $z^1 = (\underline{u}_1, e^* - \underline{u}_1)$ is obtained by having the principal pay a transfer of $(e^*)^2 + (\underline{u}_1 + \delta\underline{u}_2 - \delta e^*)/(1-\delta)$ to the agent, after which they follow the strategy profile for z^2. If the principal deviates by paying a strictly higher initial transfer, or if the agent also pays a transfer, then the deviation is ignored. If the principal deviates by transferring less, then play reverts to that specified to achieve continuation value $\underline{u}$.

The SPE strategy profiles for Model 1 are similar. To obtain $\underline{u}$, the players use the strategy profile specifying the following. Regardless of the history, the proposer should offer no externally enforced or bonus transfers. The responder should reject this and any offer that gives the responder an externally enforced transfer less than the responder's outside-option payoff (triggering disengagement), and accept any deviant offer that would give the responder an externally enforced transfer greater than the responder's outside option payoff. If the parties are engaged, the agent should exert zero effort. Finally, the principal should never make any bonus payment, regardless of the history (including whatever was promised in the negotiation phase). It is clear that players could not gain by deviating in any contingency.

To obtain z^2, the following is specified in a given period so long as there have been no deviations in the past: The proposer should offer an externally enforced transfer equal to the corresponding voluntary transfer specified in the strategy profile for Model 2, and promise a zero bonus. The responder should accept this offer, as well as any deviant offer that gives the responder an externally enforced transfer greater than the responder's outside-option payoff. Other offers should be rejected, triggering disengagement. If there are no deviations in the negotiation phase, then the agent is supposed to choose effort e^* in the production phase; if the players are engaged despite a deviation in the negotiation phase, then the agent is supposed to exert zero effort. The principal should not make any bonus payment, regardless of what was promised in the negotiation phase. If there are any deviations in the period, then continuation play starting in the next period reverts to that described above to obtain $\underline{u}$. One can verify that the players cannot gain by deviating in any contingency.

All other values in $v \in V^{\mathrm{SPE}}$ are achieved in equilibrium using appropriately prescribed first-period transfers from the principal and then conditioning the continuation as described to achieve z^2, similar to the construction for Model 2.

Finally, if $\underline{u}_1 + \underline{u}_2 > \delta/4(1-\delta)$ then $g(\underline{u}_1 + \underline{u}_2) = 0$ and no effort can be sustained in any equilibrium of either model. In both models, the parties should always offer and pay zero transfers. In Model 2 they should always disengage. If in Model 1 the proposer deviates to offer the responder an externally enforced transfer greater than the responder's outside option, then the responder should accept; otherwise the responder should reject. If the parties are engaged, the agent should exert zero effort. In each model this behavior constitutes a SPE within a single period, so repeating it unconditionally is a SPE in the repeated game. □

Observe that our equilibrium construction for Model 1 does not make use of its special features. The players do not make any payments in the bonus phase, and the monetary transfers prior to the production phase are the same regardless of whether they are externally enforced. The outside option can be triggered voluntarily rather than automatically upon the rejection of an offer. The equilibrium proposal does not depend on the identity of the proposer, and any nonzero recommended bonus schedule in the negotiation phase is to be ignored or punished rather than carried out. Moreover, our constructions for z^1 and z^2 in both models can easily be modified to satisfy the "strong optimality" refinement of Levin (2003) Corollary 1, which requires the continuation value to be on the equilibrium Pareto frontier at the start of the period after every history.[9]

So we see that the familiar relational-contracting analysis provides no basis for the prediction that the players would coordinate on any particular equilibrium. Imposing strong optimality assumes the conclusion: that the players will obtain a value on the Pareto frontier of the set V^{SPE}, if there exists an equilibrium that is on that frontier after every history. Although Model 1 allows an ultimatum offer of an externally enforced transfer and a suggested bonus, the bargaining power we might expect an ultimatum proposer to exert is easily subverted by ordinary repeated-game rewards and punishments. So something important is missing.

The point of modeling "contracting" is that, through negotiation, the parties actively discuss and agree on a prescription for behavior in their relationship. Relational contracts rely on self-enforcement, so active contracting is about coordinating on a strategy profile in the continuation game. To represent discussion about the strategy profile, we need a language suitable for the players to convey proposals for equilibrium behavior, not merely the transfers that will be paid in the current period. The simplest approach is to allow them to propose a continuation payoff to receive, and let the continuation strategy profile that delivers it be implicitly understood. We have also seen that transfers and disengagement decisions do not need to be externally enforced, so we will model them as voluntary deci-

[9] Specifically, if the principal deviates within a period, then the players coordinate on behavior that achieves z_2 from the start of the next period. Likewise, if the agent deviates within a period, then they coordinate on behavior that achieves z^1 from the next period. These punishments give the deviating party the same payoff as with continuation value $\underline{u}$. All values between z^1 and z^2 can be achieved in this fashion.

sions. This leads us to our preferred noncooperative model, which augments Model 2 by adding a cheap-talk negotiation phase to the beginning of each period.

Model 3 (Platform for Active Contracting)**.** *In each period, the agent and principal interact in the following three phases, in the order shown:*

- **Negotiation phase:** *Nature selects one of the players to be the proposer. The agent is selected with fixed probability π_1 and the principal is selected with probability π_2, where $\pi_1, \pi_2 \geq 0$ and $\pi_1 + \pi_2 = 1$. The proposer offers a continuation value $w \in \mathbb{R}^2$. The responder either accepts or rejects the offer. Both the proposer's offer and the responder's reply are cheap talk; i.e., they have no payoff consequences and do not constrain what the players may do in the continuation game.*

- **Transfer phase:** *The players simultaneously pay voluntary monetary transfers and can also burn money. Let $m = (m_1, m_2)$ denote the vector of net monetary transfers received, where $m_1 + m_2 \leq 0$.*

- **Production phase (stage game):** *The players simultaneously decide whether to disengage in the current period. If either player elects to disengage, then there is no production in the period and the players receive their outside-option payoffs given by vector $\underline{u}$. If they are engaged, then the agent selects an effort level $e \geq 0$ and the players receive stage-game payoff vector $u = (-e^2, e + e^2)$.*

Player i's payoff in the period is $(1-\delta)(m_i + u_i)$ if engaged and $(1-\delta)(m_i + \underline{u}_i)$ if disengaged, for $i = 1, 2$.

In subgame perfect equilibrium, the cheap-talk negotiation phase does not affect the set of SPE payoffs, because one can always construct an equilibrium in which cheap-talk messages are ignored. We thus obtain the following result.

Proposition 2. *The SPE payoff predictions in Model 3 are the same as in Model 1 and Model 2.*

In the next section we provide axiomatic foundations that enable the parties in Model 3 to actively negotiate over how to play the game, exerting bargaining power in the process. The resulting theory of *contractual equilibrium* yields specific, interpretable, and interesting behavior, where bargaining power matters, the Pareto frontier may be strictly worse than that of V^{SPE}, and after some histories the parties remain engaged even if they disagree.

3 Modeling active negotiation over self-enforced contracts

To model active contracting, we need to impart meaning to the parties' negotiation process. The parties need to be able to discuss how to play, and coordinate on an understanding of how to play when they reach an agreement. Model 3 provides a language that the proposer can use to make an explicit suggestion for their continuation play, by suggesting a payoff vector for them to attain. Similarly, we interpret the responder's acceptance as creating an "agreement", and the responder's rejection as generating a "disagreement". But this is not enough, since the parties could simply ignore the content of their communication. So the next step is to impose axioms on endogenous meaning, to restrict attention to equilibria in which the parties treat credible agreements as meaningful. But even this is not enough, as we show in Section 3.1, since the parties could still leverage disagreement play to punish deviant proposals.

In Section 3.2 we separately explain each of the three axioms that define contractual equilibrium, and in Section 3.3 we explain how they operate together. In particular, we show how our "theory of disagreement"—that continuation play under disagreement does not depend on how disagreement arose—interacts with the other axioms to enable players to reach meaningful agreements. In Section 3.4 we reinterpret contractual equilibrium in a hybrid model where bargaining follows the cooperative Nash bargaining solution, while play in the stage game is noncooperative.

3.1 Implementing credible agreements is not enough

In active contracting, the parties mutually understand what is being proposed, and should follow through on their agreement if it is suitably self-enforceable. To make this idea precise, one can impose one or more *agreement consistency* axioms. Each such axiom deems a set of agreements as credible, and imposes the constraint that if a credible agreement is reached, then it is carried out. Or, more precisely, an agreement consistency axiom disqualifies subgame perfect equilibria in which credible agreements are not carried out. The agreement consistency axioms differ in which agreements they consider credible.

We start by considering the strongest possible agreement consistency axiom: the parties honor any agreement to obtain a continuation value that is attainable in subgame perfect equilibrium. This is the strongest notion of what makes an agreement credible, as it ignores the potential contradictions that could arise if the players had previously agreed to play an equilibrium supported by Pareto-dominated punishments, and then had a joint incentive to deviate by negotiating out of a punishment. Our next result shows that even with this strong restriction, the same payoff set V^{SPE} is still attainable in equilibrium.[10]

Axiom 1 (Universal Agreement Consistency, or UAC)**.** *For every history at the start of a period, if in the negotiation phase the parties reach an agreement $w \in V^{\mathrm{SPE}}$, then their continuation value from the start of the transfer phase is w.*

Proposition 3. *The SPE payoff predictions of Model 3 under UAC are the same as the SPE payoff predictions of Model 1.*

See the Appendix for a proof of this proposition. We see that even such a strong agreement consistency axiom as UAC does not eliminate any possible equilibrium payoffs. As the proof shows, the problem is that even if the proposer makes a Pareto-improving offer, the responder can be rewarded for rejecting the offer, in a way that also punishes the proposer for making it. To obtain a refinement in which the parties can meaningfully discuss how to play, we need to rule out this kind of behavior. That is, we need a theory of disagreement.

3.2 Axioms of active contracting

To model active contracting, we introduce three axioms: two agreement consistency axioms that in combination are weaker than UAC, and a disagreement axiom. (We have slightly simplified the statement of each axiom compared to Miller and Watson (2013), but in a way that does not affect the characterization of contractual equilibrium.) The No-Fault Disagreement axiom embodies the idea that play under disagreement should not depend on how disagreement occurred. Disagreements can arise in a variety of ways: either player may be selected as proposer; then the proposer may make a deviant offer and the responder rejects it, or the proposer may make the equilibrium offer and the responder rejects it. The axiom also embodies the idea that transfers should not be made if there is no agreement.[11]

Axiom 2 (No-Fault Disagreement, or NFD)**.** *For every history at the start of a period t, there exists a disagreement value $\underline{w} \in V^{\mathrm{SPE}}$ such that if the responder rejects the proposer's offer, then the continuation value from the start of the production phase is $\underline{w}$, regardless of the identity of the proposer, the content of the offer, and what monetary transfers were made in period t.*

[10]This result strengthens Theorem $1'$ in Miller and Watson (2013), which imposes a weaker set of agreement consistency axioms.

[11]The intent here is to model a situation in which there is no external enforcement. Watson, Miller, and Olsen (2020) allows for long-term externally enforced contracts, in which case a legal contract signed in a prior period would be externally enforced even under disagreement in the current period.

The axiom still allows that the disagreement value $\underline{w}$ may differ depending on the history. For instance, the disagreement value may favor the agent if the agent exerted effort e^* in the prior period, but may favor the principal if the agent exerted deviant effort. The axiom also implies that under disagreement neither player will be willing to make a voluntary transfer, since the continuation value after the transfer phase does not depend on what transfers were made. By itself, No-Fault Disagreement does not materially affect the set of equilibrium payoffs, because it does not constrain play under agreement; this is shown in Theorem 2 of Miller and Watson (2013). We will shortly see that when NFD is combined with the agreement consistency axioms, the agreement the parties will reach in a given period is heavily influenced by what they would get if they disagreed. With the flexibility to condition their future disagreement values on the history of play, the parties can implicitly condition their future agreements on the history as well, and thereby provide incentives for them to conform to their current agreement.

The agreement consistency axioms enable the parties to reach agreement on credible continuation play. Internal Agreement Consistency (IAC) deems an agreement w credible if it is "supported within the current equilibrium." For a given equilibrium there is a set W containing every value that can be obtained starting from the production phase, at some history. An offer w is *supported within the current equilibrium* if it can be obtained by first making an arbitrary transfer and then continuing from the production phase with a value in W. We provide two examples to illustrate what W may look like.

Example 1. *Consider an equilibrium in which the players are supposed to obtain a continuation value of $\underline{u}$ after every history. In this case $W = \{\underline{u}\}$. While the players could feasibly obtain other values by first exchanging a transfer and then continuing with this behavior, neither player would ever be willing to transfer a non-zero amount.*

Example 2. *In the proof of Proposition 1 we described a subgame perfect equilibrium in which at some histories to the production phase the agent exerts effort e^* and then obtains z^2 in the following period, so W contains the value $(1-\delta)(-(e^*)^2, 0) + \delta z^2$. At other histories (following a deviant transfer by the principal) the parties disengage and then continue with $\underline{u}$, so W also contains the value $(1-\delta)\underline{u} + \delta\underline{u} = \underline{u}$. Other continuation values were obtained by first exchanging a transfer and then continuing with one of these values. Thus $W = \left\{(1-\delta)(-(e^*)^2, 0) + \delta z^2, \underline{u}\right\}$.*

IAC allows the parties to discuss whether to switch to an alternative history using the continuation values already present in their current equilibrium. Under IAC, if the parties find themselves negotiating at a history at which they are supposed to endure a Pareto-dominated punishment, the proposer may instead propose switching to an alternative history at which they would receive a Pareto-superior value, and the responder may accept, upon which they should continue as agreed. So the IAC axiom enables players to renegotiate out of Pareto-dominated punishments, but does not require them to do so.

Axiom 3 (Internal Agreement Consistency, or IAC)**.** *For a given SPE, let W be the set of equilibrium continuation values from the start of the production phase, across all histories. The equilibrium satisfies IAC if, for any history to a given period, if*

1. *the offer w satisfies $w = (1-\delta)m + w'$ for a transfer m satisfying $m_1 + m_2 \leq 0$ and a continuation value $w' \in W$,*

2. *and this offer is accepted,*

then the equilibrium continuation value from the start of the transfer phase is w.

The second agreement consistency axiom is Pareto External Agreement Consistency (PEAC), which deems an agreement w credible if it is Pareto optimal among values that are supported within any equilibrium satisfying IAC and NFD.

Axiom 4 (Pareto External Agreement Consistency, or PEAC)**.** *Let $\overline{W}$ be the set of equilibrium continuation values that maximize the joint payoff from the start of the production phase, across all histories of all subgame perfect equilibria satisfying IAC and NFD. An equilibrium satisfies PEAC if, for any history to a given period, if*

1. *the offer w satisfies $w = (1-\delta)m + w'$ for a transfer m satisfying $m_1 + m_2 \leq 0$ and a continuation value $w' \in \overline{W}$,*

2. *and this offer is accepted,*

then the equilibrium continuation value from the start of the transfer phase is w.

PEAC gives the parties the ability to discuss which equilibrium to play. They can select any equilibrium that is optimal among those that accord with how they expect their future negotiations to play out—namely, any equilibrium that satisfies IAC and NFD, and will not itself be renegotiated to an even better such equilibrium.[12]

[12]This is why PEAC considers only Pareto optimal agreements credible. Contradictions would arise if a Pareto-dominated agreement were also considered credible, since there would be histories at which the players would be expected to carry out their Pareto-dominated agreement and yet also have the opportunity to renegotiate out of it to a Pareto-superior equilibrium.

Remark 1. *The renegotiation-proofness literature has also recognized that Pareto-dominated punishments may not be credible if they are vulnerable to renegotiation. The approach in that literature has been to disqualify equilibria in which renegotiation could be jointly profitable, rather than to model renegotiation. By ruling out the possibility of Pareto-dominated continuation values at all histories, renegotiation proofness rules out the possibility of disagreement, which puts it at odds with both the cooperative and non-cooperative approaches to bargaining. Despite the difference in approach, elements of the renegotiation approach relate to our axioms with regard to which continuation values they define as credible agreements. In particular, our IAC axiom relates to internal or "weak" renegotiation-proofness notions studied by Bernheim and Ray (1989), Farrell and Maskin (1989), and Ray (1994), because it deems credible the continuation values already contained in the equilibrium. Similarly, our PEAC axiom relates to external or "strong" renegotiation-proofness notions (Bernheim and Ray 1989; Farrell and Maskin 1989; Asheim 1991), because it deems credible the continuation values of certain other equilibria, outside the current equilibrium.*

Remark 2. *The axiomatic underpinnings of contractual equilibrium are modular, in the sense that NFD could be replaced with a different theory of disagreement, and the agreement consistency axioms IAC and PEAC could be replaced with different notions of agreement consistency. We have shown that the theory of disagreement should not be too permissive (Proposition 3),*[13] *but it also should not be too restrictive. If play under disagreement were anchored to single stage game outcome regardless of the history, then under reasonable agreement consistency axioms the parties would always negotiate to the same agreement, and there would be no way to provide incentives in the production phase.*[14] *One could translate the various approaches in the renegotiation-proofness literature into an alternative set of agreement consistency axioms.*[15] *It would also be interesting to incorporate more nuance into the negotiation protocol, to allow for negotiating costs (Blume 1994),*

[13] Safronov and Strulovici (2018) study a more permissive theory of disagreement that allows deviant proposals to be punished, and find that while some low payoff vectors are ruled out, typically both Pareto efficient and inefficient payoff vectors are attainable under their refinement.

[14] Ramey and Watson (2002) develop a theory of this sort to model how an external and costly dispute-resolution system can help enforce contracts. Klimenko, Ramey, and Watson (2008) is an application to international trade. Productive incentives could also be provided under the assumption, perhaps unrealistic but relatively common in the macro-labor literature, that the relationship is severed automatically in the event that the agent deviates in the production phase. Models of this type are evaluated by den Haan, Ramey, and Watson (1999, 2000a); worker bargaining power is shown to provide a better foundation for "efficiency wages" than does a firm's incentives regarding contract offers.

[15] Bergin and MacLeod (1993) map the relationships among a wide variety of axioms in the early renegotiation-proofness literature. Each notion delineates what is a credible renegotiation target, and can be reinterpreted as an agreement consistency axiom. Abreu, Pearce, and Stacchetti (1993) developed an approach in which minimally sufficient punishments are considered credible, also explained in Pearce and Stacchetti (2022).

infrequent renegotiation (Goldlücke and Kranz 2029), or an explicit account of the external enforcement process.

3.3 Contractual equilibrium

The NFD, IAC, and PEAC axioms together define contractual equilibrium.

Definition 1. *A SPE of Model 3 is a* ***contractual equilibrium*** *if it satisfies NFD, IAC, and PEAC.*

A contractual equilibrium has a specific structure, which we can deduce directly from the axioms. Consider the negotiation phase at some history h. By NFD, there is a specific disagreement value $\underline{w}$ that the parties know will arise if they fail to agree. By IAC, they can agree to obtain any payoff vector that is attained at any other history starting from the production phase in the same equilibrium, and augment it with an up-front transfer m. So if there is any production phase history h' at which the continuation value w' adds up to a higher joint payoff than that of their current disagreement point, then the proposer (player i) can offer a continuation value of $w = m + w'$. If w Pareto dominates the disagreement value $\underline{w}$, then the responder (player j) will be willing to accept. Since the proposer makes an ultimatum offer, in equilibrium the proposer will extract all the surplus by offering a continuation payoff that gives the responder $\underline{w}_j$ and gives the proposer $w'_i + w'_j - \underline{w}_j$.

Among all equilibria with this structure, by PEAC the proposer will optimally choose as w' a $\bar{w} \in \overline{W}$ that maximizes the joint payoff among all equilibria satisfying IAC and NFD. So in any given period, under these axioms there is a straightforward bargaining process arising from a standard Ultimatum Game analysis: There is a disagreement point, which may be below the Pareto frontier of what can be attained under IAC and NFD; the proposer chooses a continuation value on the Pareto frontier, augmented with a transfer that gives the responder his or her disagreement payoff and gives the rest to the proposer.

From the perspective of the start of the period, before the proposer is selected, the players' expected continuation value will depend on π_1 and π_2, the proposer selection probabilities. Specifically, in expectation they split the bargaining surplus in proportion to their probabilities of being selected to propose, which we can now interpret as bargaining weights.

Since the proposer offers a continuation value $\bar{w}$ from the production phase that is not sensitive to the history, it follows that the payoff set under agreement lies in a straight line of slope -1 through $\bar{w}$. The *contractual equilibrium value (CEV)* set W^* is the convex hull of the expected payoffs from the start of the period, across all histories. We call the joint value $\bar{w}_1 + \bar{w}_2$ the *level* of the CEV set, written as $\mathrm{level}(W^*)$.

In the context of our simple principal-agent setting, if the parties are sufficiently patient then $\text{level}(W^*)$ is attained by exerting a particular level of effort $\bar{e}$, which depends on δ, whenever they agree. Our specific functional form assumptions are calibrated so that $\text{level}(W^*) = \bar{e}$. If the parties are not sufficiently patient, they will not be able to sustain any effort in equilibrium, and should disengage each period. In what follows we focus on the case in which they are sufficiently patient, and we will note what constitutes sufficient patience.

The parties' joint value $\bar{e}$ is constrained by the strength of incentives available in the equilibrium. The incentive for the agent to engage and exert effort in the production phase of the current period is provided by transitioning to a high continuation value for the agent within W^* if she exerts the right effort, but to her worst continuation value in W^*, which we call z^1, if she deviates. Depending on whether further incentives are needed to induce the principal to engage, the best continuation value for the agent, which we call z^2, may be used either to reward the agent or to punish the principal. (The superscript indicates which player is being punished, as in the SPE sets constructed earlier.) The larger is the difference between z^2 and z^1, the higher-powered are the incentives.[16] We call this difference the *span* of W^*, computed as $\text{Span}(W^*) = z_1^2 - z_1^1$. The foregoing chain of logic indicates that the level of the CEV set is closely linked with its span, where a longer span enables a higher level.

Consider a history to the start of a period at which z^i is the continuation value. By NFD, there is a particular disagreement value $\underline{w}^i$ to anchor their negotiation process. Given $\underline{w}^i$, z^i must deliver joint payoff $\text{Level}(W^*)$ and split the surplus in proportion to the players' bargaining weights. Specifically,

$$z^i = \underline{w}^i + \pi \cdot \big(\text{Level}(W^*) - \underline{w}_1^i - \underline{w}_2^i\big), \tag{1}$$

where $\pi \equiv (\pi_1, \pi_2)$. Accordingly, to characterize the CEV set it suffices to characterize $\underline{w}^1$, $\underline{w}^2$, and $\text{Level}(W^*)$.

This characterization is accomplished for general repeated games by Theorem 5 in Miller and Watson (2013). In Proposition 4, in the Appendix, we specialize the characterization to Model 3. Proposition 4 provides a recipe for computing the CEV set, by simultaneously solving three inter-related optimization problems. Each optimization problem involves finding a stage game action profile and a transfer of continuation value such

[16] The parties are deterred from disengaging similarly: whomever deviates by disengaging is punished by their worst continuation value. If both simultaneously deviate, it doesn't matter how they continue, so for concreteness say they continue with z^2.

that the action profile is enforced.

3.4 Modeling a relationship as a hybrid game

Contractual equilibrium is what results when axioms on endogenous meaning are applied to a non-cooperative repeated game model with a cheap-talk bargaining phase at the start of each period. But it turns out that the equilibria that satisfy these axioms are outcome-equivalent (in terms of expected payoffs) to what arises in a hybrid model that replaces the non-cooperative cheap-talk bargaining phase in each period with a cooperative bargaining model. That is, the bargaining phase is modeled as a joint decision, reached according to the Nash bargaining solution (Nash 1950).

Model 4 (Hybrid)**.** *In each period, the agent and principal interact in the following two phases, in the order shown:*

- **Negotiation phase:** *The players cooperatively bargain over a continuation plan and a balanced monetary transfer $m = (m_1, m_2)$, where $m_1 + m_2 = 0$. The vector of exogenous bargaining weights is $\pi = (\pi_1, \pi_2)$, where $\pi_1, \pi_2 \geq 0$ and $\pi_1 + \pi_2 = 1$. If they fail to agree, they continue without a transfer and follow the plan they most recently agreed on in a prior period, or the default plan if they have never agreed. If they agree, the transfer is instantiated immediately.*

- **Production phase (stage game):** *The players simultaneously decide whether to disengage in the current period. If either player elects to disengage, then there is no production in the period and the players receive their outside-option payoffs given by vector $\underline{u}$. If they are engaged, the agent selects an effort level $e \geq 0$ and the players receive stage-game payoff vector $u = (-e^2, e + e^2)$.*

In the period, each player i obtains $(1 - \delta)(m_i + u_i)$ if engaged and $(1 - \delta)(m_i + \underline{u}_i)$ if disengaged.

The negotiation phase is resolved according to the Nash bargaining solution, where the bargaining set is the set of equilibrium continuation values and the disagreement point is the value of disagreement play. The production phase is resolved according to recursive Nash equilibrium, where the stage game is augmented with conditional expected continuation values. Any equilibrium value set consistent with these conditions will correspond to an equilibrium of Model 3 under the IAC and NFD axioms. The equilibrium value set with the highest level is the CEV set. We have outlined the approach only informally here; the formal details are laid out in Watson, Miller, and Olsen (2020). While the hybrid model is not widely used in the relational contracts literature, we recommend it as a simpler reduced form of the non-cooperative model when employing contractual equilibrium. Conveniently, the CEV set is characterized by a recursive formulation along the lines of Abreu, Pearce, and Stacchetti (1990); see Watson (2021) for an overview.

4 Contractual equilibrium in the principal-agent model

To illustrate the interesting properties that arise in contractual equilibrium, in this section we describe contractual equilibrium for the special case of $\underline{u} = (0, 0)$. The mathematical details are reserved to the appendix. In this special case, we can ignore the decision over whether to disengage, since disengagement yields the same payoff as engagement with zero effort.[17] Remarkably, the axioms suffice to identify behavior that is essentially unique, in the sense that there is a unique strategy profile that the parties will play in any continuation game following their first agreement and first transfer. The behavior is also interesting and intuitive. An important property is that the level of effort sustainable in equilibrium is $\pi_1\delta/(1 - \delta)$, which is proportional to the agent's bargaining weight.

Effort $\overline{e}$ that gives the level of the CEV set is increasing in the span of the CEV set, because the agent's incentive to exert effort is maximized by using the full span of continuation values to reward and punish her. Specifically, play of $\overline{e}$ is rewarded with continuation value z^2 from the next period, whereas any deviation would be punished with continuation value z^1. Thus, $\overline{e}$ is the largest effort e satisfying $(1 - \delta)e^2 \leq \delta \operatorname{Span}(W^*) = \delta(z_1^2 - z_1^1)$, implying $\overline{e} = \sqrt{(z_1^2 - z_1^1)\delta/(1 - \delta)}$. Points on the line segment between z^1 and z^2 are achieved by having player 1 select $\overline{e}$ with the rewards and punishments just described, preceded by a transfer in the current period to shift utility from one player to the other.

Next we characterize the endpoints of the CEV set, z^1 and z^2. The key to the analysis—in fact, the key to understanding the contractual equilibrium construction overall—is to recognize how the disagreement value in one period depends on the agent's effort in the prior period. Recall that z^1 is the value of negotiation from disagreement point $\underline{w}^1$, and likewise z^2 is the value of negotiation from disagreement point $\underline{w}^2$. Thus $\underline{w}^1$ should be the value that most favors player 2 and $\underline{w}^2$ should be the value that most favors player 1, among the possible continuation values under disagreement.

Suppose the agent exerted the desired effort in the prior period, and if the parties fall into disagreement in the current period then they continue with the value $\underline{w}^2$ at the start

[17] This special case avoids needing to deter the principal from disengaging when under disagreement. A full analysis with arbitrary $\underline{u}$ is sufficiently rich as to be outside the scope of this paper.

of the production phase. To make this continuation value most favorable to the agent, the equilibrium specifies that the agent should not exert any effort, and starting in the next period the agent gets her highest equilibrium payoff. In this scenario, the disagreement value for the negotiation phase is $\underline{w}^2 = (1-\delta)(0,0) + \delta z^2$. At the start of the negotiation phase, the parties expect to bargain up to the equilibrium-path effort, splitting the additional surplus generated in proportion to their bargaining weights. That is, $z^2 = \underline{w}^2 + \pi\left(\mathrm{Level}(W^*) - \underline{w}_1^2 - \underline{w}_2^2\right)$, which implies $z^2 = \pi\,\mathrm{Level}(W^*)$. To see this, combine the two equations and use $z_1^2 + z_2^2 = \mathrm{Level}(W^*)$, which follows from the second equation, to simplify.

Next suppose that the agent deviated in the production phase of the prior period. The disagreement value in this case, $\underline{w}^1$, entails the agent selecting $\overline{e}$ in current period and being rewarded with continuation value z^2 from the next period, whereas she would be punished for any deviation by reverting to continuation value z^1 from the next period. We therefore have $\underline{w}^1 = (1-\delta)\left(-\overline{e}^2, \overline{e} + \overline{e}^2\right) + \delta z^2$. Note that there is no negotiation surplus with this disagreement point, and so $\underline{w}^1 = z^1$. The agent's part, $\underline{w}_1^1$, must equal $(1-\delta)0 + \delta z_1^1$ due to the agent's binding incentive condition, implying $z_1^1 = \delta z_1^1$ and thus $z_1^1 = 0$. As a result, the same behavior in the production phase and continuation play is called for under both agreement and disagreement. This conclusion stands in stark contrast to the assumption in Levin (2003) that disagreement immediately causes disengagement.

The preceding analysis shows that effort $\overline{e}$ and the endpoints of the CEV set satisfy $\overline{e} = \sqrt{(z_1^2 - z_1^1)\delta/(1-\delta)}$, $z_1^2 = \pi_1 \overline{e}$, and $z_1^1 = 0$. These equations imply $\overline{e} = \pi_1\delta/(1-\delta)$, and therefore $z^1 = (0, \pi_1\delta/(1-\delta))$ and $z^2 = \pi \cdot \pi_1\delta/(1-\delta)$. The CEV set is pictured in Figure 3. The endpoints z^1 and z^2 are shown along with their corresponding disagreement points $\underline{w}^1$ and $\underline{w}^2$.

In equilibrium, if the agent exerts the desired effort, then starting the in next period she gets a π_1 share of the surplus (at z^2), whereas if she deviates then she gets an expected payoff of zero (at z^1). This also stands in contrast to the Levin (2003) conclusion that the entire relationship surplus can be used to motivate the agent. There are two keys to understanding the distinction. First is that because the standard approach, exemplified by Levin (2003), does not model active contracting, bargaining power does not play a role. Instead, the standard approach gives none or all of the surplus to the agent depending on whether the agent deviated. In contrast, by modeling active contracting, contractual equilibrium incorporates bargaining power, so each party receives their share of the surplus at every successful negotiation. Second, Levin (2003) assumes that the outside option is triggered automatically in case of disagreement. In the Levin (2003) analysis this doesn't matter, as we showed

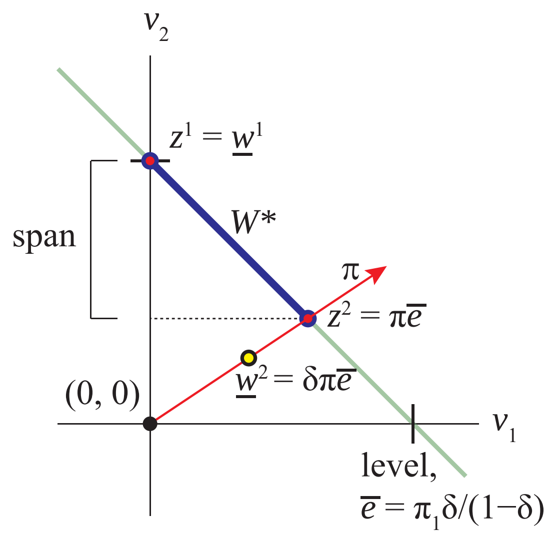

FIGURE 3. Principal-agent CEV set.

in Proposition 1. But if that assumption were imposed here—as an alternative theory of disagreement in place of the NFD axiom—then no effort could be sustained at all! Positive effort is sustainable in contractual equilibrium only if the behavior under disagreement differs depending on whether the agent is being rewarded or punished.

Remark 3. *As the principal-agent example illustrates, when in a contractual equilibrium the players coordinate on any particular continuation value w from the start of a period, it is effectively the result of coordinating on an associated disagreement value $\underline{w}$ from which negotiation would lead to w. One might ask whether the theory necessarily predicts that renegotiation will occur in every period, even on the equilibrium path, which seems at odds with the workings of real contractual relationships. In fact, a minor enrichment of the model supports an equilibrium construction without such renegotiation. Specifically, adding to Model 3 or Model 4 a voluntary transfer phase at the end of the period (what was called the bonus phase in Model 1) does not change the CEV set but allows for a convenient equilibrium construction in which active renegotiation doesn't need to occur unless there is a deviation from prescribed voluntary transfers at the end of a period.*[18] *Something similar*

[18] For example, consider an equilibrium that, after a particular history, specifies coordination on disagreement value $\underline{w}$ that the players would renegotiate from to achieve $w = \hat{w} + (m_1, m_2)$, where (m_1, m_2) is the negotiated transfer and $\hat{w}$ is the continuation value from the production phase that the players negotiate to obtain. This equilibrium can be modified so that a transfer of $\delta(m_1, m_2)$ occurs at the end of the previous period, following which the players would coordinate to achieve $\hat{w}$ without renegotiation (since $\hat{w}$ is a feasible

can be done for settings with nontrivial external enforcement, as Watson, Miller, and Olsen (2020) explain.[19]

5 Conclusion

We have contrasted the contractual-equilibrium approach, based on axioms that address how players interpret bargaining outcomes that would otherwise be mere cheap talk, with the standard approach of selecting an equilibrium on the Pareto frontier of the SPE payoff set. Under the standard approach, various modeling features—such as ultimatum contract offers, promised bonuses, enforceable wages, and automatic disengagement upon disagreement—have no effect on the set of attainable equilibrium values, leaving the equilibrium selection problem just as much in the hands of the analyst as they would be without these features. Contractual equilibrium, in contrast, identifies essentially unique equilibrium behavior in many settings, including the canonical principal-agent game that we use as an example.

Under contractual equilibrium, we find that bargaining power matters, as does play under disagreement. In the principal-agent game, the incentive power given to the agent is bounded by the agent's bargaining share of the relationship surplus. Incentives are supported by playing differently under disagreement depending on whether the agent is being rewarded or punished, so that the agreements formed relative to these disagreement outcomes feature high and low transfers to the agent, respectively.

The theory of negotiation underlying the contractual equilibrium concept has a familiar characterization as the generalized Nash bargaining solution applied to a bargaining set and disagreement point. The parties share the negotiation surplus in proportion to their fixed bargaining weights π, and the disagreement point is any history-dependent continuation value that can be achieved without an immediate transfer. Bargaining power derives from the bargaining weights, along with the endogenous disagreement point. This characterization brings the relational-contracting literature closer to other lines of scholarship that take bargaining power seriously, including the large literature on hold-up problems (see, for instance, Grout 1984), property rights and "incomplete contracts" (Grossman and Hart 1986,

disagreement value and is already on the frontier of the CEV set). Noting that $m_1 = -m_2$, if the player expected to make the transfer fails to do so, then the players coordinate on disagreement value $\underline{w}$ and expect to renegotiate as in the original equilibrium.

[19]The modification with external enforcement is a little more complicated because it involves sending messages to the external enforcer. Incidentally, in settings with more than two players and/or imperfect joint monitoring, including an end-of-period transfer phase augments the production technology is a way that can expand what can be attained in contractual equilibrium.

Hart and Moore 1990, and others), and macro-labor market dynamics (such as Mortensen and Pissarides 1994 on job creation and destruction, den Haan, Ramey, and Watson 2000b on propagation of shocks, and den Haan, Ramey, and Watson 2003 on liquidity flows). Models in these categories similarly take the bargaining weights as exogenous and fixed, typically as parameters to either calibrate or estimate (as in Svejnar 1986). Some models with moral hazard assume that deviations in the production phase directly cause separation, displacing the kind of self-enforced punishment paths that arise in contractual equilibrium.

While our main goal is to understand how relational contracts operate, there is a normative flavor to our approach. For relational contracts to succeed, the parties should recognize and account for the possibility of disagreement, and explicitly agree on their plan for what should happen in case of future disagreements. The No-Fault Disagreement axiom further suggests that players should not hold each other liable for rejected deviant offers. However, this notion cuts two ways: it rules out the bad outcomes that could arise from discouraging innovative offers, but it can also rule out superior outcomes that could potentially be supported by carefully calibrating the endogenous responses to rejected deviant offers. Our view is that a relationship will be more resilient if negotiations are founded on a realistic appraisal of relative bargaining power, rather than on the prospect that the parties can be jointly relied on to reward one of them for rejecting a tempting out-of-equilibrium proposal. In Gjertsen, Groves, Miller, Niesten, Squires, and Watson (2021) we took a more explicitly normative approach, using contractual equilibrium as a lens with which to interpret successes and failures in several case studies of long-run relationships.

What is the source of bargaining power remains a deep open question. In our noncooperative modeling, formally bargaining weights represent power to set the agenda by making proposals. So if an institution governs the bargaining protocol, then it can allocate bargaining weights by design. It is important to note that bargaining weights do not arise from outside options, control rights, or asset ownership—these affect the actions available to the agents in case there is a disagreement (e.g., the disengagement choice in the stage game of Model 3), but do not affect the division of surplus relative to the disagreement point. In our hybrid model, bargaining weights are simply the share of surplus that each party is understood to be entitled to.[20]

In constructing a contractual equilibrium for the principal-agent game, we limited our

[20] An alternative to founding bargaining weights on the bargaining protocol is that that they may be driven by social norms. Ghosh and Ray (2022) and MacLeod and Malcomson (2021) discuss how social norms may arise. Under the MacLeod and Malcomson (1998) approach, the social norm would allocate bargaining power so as to maximize the surplus of the relationship, which in the principal-agent example of Section 4 would mean allocating all bargaining power to the agent.

attention to the case in which the outside option yields the same payoff vector that would arise if the agent exerted zero effort. Extending the analysis to allow for higher outside options is not particularly difficult under the standard approach (as detailed in Levin 2003), but under contractual equilibrium there are new subtleties to consider. While these are outside the scope of this paper, which is intended to be heuristic, they are quite interesting. Briefly, the issue is that while an attractive outside option will constrain what is enforceable under disagreement if the parties do engage, it also provides the additional possibility of disengagement. To determine whether it is optimal to disengage under disagreement when rewarding or punishing the agent will require a detailed comparison that will naturally depend on the nature of the outside option. This is an interesting question for future research, since understanding the effect of outside options can help us understand how relational contracts between two parties are influenced by the broader environment in which such relationships are formed and maintained.

Contractual equilibrium can, of course, be applied to much broader class of economic settings than the complete-information principal-agent model we analyzed here. Both of the more general models considered by Levin (2003)—one with imperfect monitoring, and one with private information—fall within the class of imperfect public monitoring models considered by Miller and Watson (2013). Endogenous monitoring (as in Gjertsen, Groves, Miller, Niesten, Squires, and Watson 2021) and team production are two additional areas of application, among many others. Miller and Watson (2013) also showed how to handle heterogeneous discount factors, along with a simple way to address relational contracts among more than two players. Watson (2013) formalizes contractual equilibrium for finite-period relationships. More recently, Watson, Miller, and Olsen (2020) extend the scope of contractual equilibrium to allow for external enforcement of arbitrary long-term contracts. Kostadinov (2021) further extends the scope (including long-term contracts) to allow the agent to be risk averse, so that monetary transfers are not synonymous with transfer of utility.[21] With external enforcement of long-term contracts, the contract signed in one period determines the environment in which the parties will interact if they disagree in the following period and fail to renegotiate the contract.

Looking forward, contractual equilibrium can be applied to understanding relational contracts in non-stationary environments, such as with accumulation of capital, innovation, limited liability for the agent, or financing constraints for the principal. Such environments naturally feature equilibrium multiplicity, and contractual equilibrium provides a disciplined equilibrium selection. It is also important to understand contractual relation-

[21] Kostadinov's contributions are also explained in Pearce and Stacchetti (2022).

ships in their larger context, whether that be a matching market where parties can search for new relationships upon separation (see Fahn 2017 for a modern analysis of efficiency wages along this line), or a network of overlapping relationships. In some cases a large player may be able to design or at least influence the environment, such as a platform that connects agents and principals to form matches. Bernstein (2016) explains how a large firm with many suppliers can overcome its own hold-up incentives by connecting its suppliers to each other, facilitating collective punishments—even while each relationship between the firm and a given supplier is governed by a bespoke contract. In such an environment, the disagreement point in a given bilateral relationship could depend on the state of the multilateral relationship.

References

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 infinite horizon games. *Games and Economic Behavior*, 3(3):278–294, August 1991.

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.

Lisa Bernstein. Beyond Relational Contracts: Social Capital and Network Governance in Procurement Contracts. *Journal of Legal Analysis*, 7(2):561–621, 02 2016. ISSN 2161-7201. doi: 10.1093/jla/law001. URL https://doi.org/10.1093/jla/law001.

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

Wouter J. den Haan, Garey Ramey, and Joel Watson. Contract-theoretic approaches to wages and displacement. *Federal Reserve Bank of St. Louis Review*, (May):55–68, 1999.

Wouter J. den Haan, Garey Ramey, and Joel Watson. Job destruction and the experiences of displaced workers. In *Carnegie-Rochester Conference Series on Public Policy*, volume 52, pages 87–128. Elsevier, 2000a.

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

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

Matthias Fahn. Minimum wages and relational contracts. *The Journal of Law, Economics, and Organization*, 33(2):301–331, 2017.

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

Parikshit Ghosh and Debraj Ray. The social equilibrium of relational arrangements. Working pape, May 2022.

Heidi Gjertsen, Theodore Groves, David A. Miller, Eduard Niesten, Dale Squires, and Joel Watson. Conservation agreements: Relational contracts with endogenous monitoring. *Journal of Law, Economics, and Organization*, 37(1):1–40, March 2021.

Susanne Goldlücke and Sebastian Kranz. Reconciliating relational contracting and hold-up: A model of repeated negotiations. Working paper, December 2029. URL https://ssrn.com/abstract=2210603.

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.

Paul A. Grout. Investment and wages in the absence of binding contracts: A nash bargaining approach. *Econometrica*, 52(2):449–460, 1984.

Marina Halac. Relational contracts and the value of relationships. *American Economic Review*, 102 (2):750–779, April 2012.

Marina Halac. Revealing the value of relationships. Working paper, March 2022.

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

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.

Rumen Kostadinov. Renegotiation of long-term contracts as part of an implicit agreement. *Theoretical Economics*, 16(4):1471–1512, November 2021.

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

W. Bentley MacLeod and James M. Malcomson. Motivation and markets. *American Economic Review*, 88(3):388–411, June 1998.

W. Bentley MacLeod and James M. Malcomson. Implicit contracts, incentive compatibility, and involuntary unemployment: Thirty years on. Working paper, December 2021.

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

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.

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

David Pearce and Ennio Stacchetti. Relational contracts: Reputation and renegotiation. Working paper, March 2022.

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

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

Mikhail Safronov and Bruno Strulovici. Contestable norms. Working paper, December 2018.

Jan Svejnar. Bargaining power, fear of disagreement, and wage settlements: Theory and evidence from US industry. *Econometrica: Journal of the Econometric Society*, pages 1055–1078, 1986.

Joel Watson. Contract and game theory: Basic concepts for settings with finite horizons. *Games*, 4: 457–496, 2013.

Joel Watson. Theoretical foundations of relational incentive contracts. *Annual Review of Economics*, 13:631–659, 2021.

Joel Watson, David A. Miller, and Trond E. Olsen. Relational contracting, negotiation, and external enforcement. *American Economic Review*, 110(7):2153–2197, July 2020.

A Appendix

A.1 Proof not shown in the main text

This appendix section proves Proposition 3, which states that the SPE payoff predictions of Model 3 under UAC are the same as the SPE payoff predictions of Model 1.

Proof. Given any $v \in V^{\text{SPE}}$, we construct a subgame perfect equilibrium that delivers value v while satisfying UAC. First, we describe the equilibrium path behavior. In period 1, the proposer is supposed to offer v, and the responder is supposed to accept. In any period t, suppose $w \in V^{\text{SPE}}$ is the agreement reached in the negotiation phase; it becomes the "standing agreement" for subsequent phases. The parties should exchange monetary transfers m that satisfy

$$w = (1-\delta)\Big(m + \big(-(e^*)^2, e^* + (e^*)^2\big)\Big) + \delta z^2,$$

neither party should disengage, and the agent should exert effort e^*. In period $t+1$ the proposer should offer z^2. The responder should accept. Then they continue their play, with z^2 as the standing agreement.

So far we have partially described a strategy profile that satisfies UAC on the equilibrium path, and delivers value v. Next we complete the description of the strategy profile and explain why it is a subgame perfect equilibrium. Given the outcome of the random proposer selection in a given period, label the proposer as player i and the responder as player j.

The key step addresses what happens if the proposer deviates to offer $w' \in V^{\text{SPE}}$ when they were supposed to offer w in the negotiation phase of any period t. Then the responder should reject, and the parties should behave as if the standing agreement were instead z^i. The reward of $z^i_j \geq w'_j$ makes the responder willing to reject the credible offer, and the penalty of $z^i_i \leq w_i$ discourages the proposer from offering it. Because the responder rejects the offer, UAC does not constrain their continuation play. (For any other deviant offer $w' \notin V^{\text{SPE}}$, the responder may simply accept, and the parties should behave as if the standing agreement were instead w; this does not violate UAC, since it does not deem an agreement $w' \notin V^{\text{SPE}}$ credible.)

If the responder deviates by rejecting an offer that was supposed to be accepted, then they can simply behave as if the offer was accepted, since UAC does not apply to behavior under disagreement. Thus the responder is willing to accept offers that are supposed to be made in equilibrium. (If the responder deviates by accepting an offer w that was supposed to be rejected, our construction implies that the offer satisfies $w \in V^{\text{SPE}}$, so by UAC they must continue with w as the standing agreement.)

If either player deviates from the specified monetary transfer in period t, then both parties should disengage, and in period $t+1$ the proposer should offer $\underline{u}$ and the responder should accept. If either party deviates from the specified engagement decision in period t, then in period $t+1$ the proposer should offer $\underline{u}$, and if they are engaged the agent should exert zero effort. Since the specified monetary transfer is always in service of obtaining some continuation value $w \in V^{\text{SPE}}$, the threat of disengagement followed by $\underline{u}$ suffices to discourage deviant transfer and engagement decisions.

If the agent deviates in the production phase in period t, then in period $t+1$ the proposer should offer z^1, and the responder should accept. Then they continue their play, with z^1 as the standing agreement. The reward of getting z_1^2 rather than z_1^1 makes the agent willing to exert effort e^*, as established in Proposition 1. $\square$

A.2 Characterization of the CEV set

This appendix section provides details for how to compute the CEV set, by specializing Theorem 5 in Miller and Watson (2013) here to Model 3. For simplicity we restrict attention to pure strategies in the production phase, which in this particular model is without loss of generality.

To account for both engagement and effort decisions in the same stage game payoff function, let $f \in \{0,1\}^2$ be the pair of indicators for whether the agent and principal, respectively, chose to engage; let $F \in \{0,1\}$ be the indicator for whether they are engaged. Then (with some abuse of notation) we define $u(e,f) = Fu(e) + (1-F)\underline{u}$. The parties' interaction starting in the production phase and then continuing with a value from the CEV set next period can be summarized by augmenting their production phase payoffs in the current period with a transfer of continuation value $\eta : \mathbb{R}_+ \times \{0,1\}^2 \to [0, \text{Span}(W^*)]$ from principal to agent, so for purposes of verifying incentives their payoffs can be normalized to $(1-\delta)u(e,F) + \delta(\eta(e,f), -\eta(e,f))$. Then we say the action profile $(e,f) \in \mathbb{R}_+ \times \{0,1\}^2$ is *incentive compatible (IC) under* η if[22]

$$(1-\delta)u_1(e) + \delta\eta_i(e,(1,1)) \geq (1-\delta)u_1(0) + \delta\eta(0,(1,1)); \tag{2}$$

[22]Without loss of generality, we restrict attention to strategy profiles that call upon both players to take the same engagement/disengagement action. Without loss of generality, our incentive compatibility condition focuses on deterring the agent from exerting zero effort: If η suffices to deter deviating from e to zero effort, then setting $\eta(e') = \eta(0)$ for all $e' \neq e$ deters all other effort deviations without violating any of the constraints. We ignore the incentive conditions for both players to disengage when $F = 0$, since if both are called on to disengage, then each is indifferent over whether to do so.

and we say it is *individually rational (IR) under* η if $f = (0,0)$, or if $f = (1,1)$ and

$$(1-\delta)u_1(e,(1,1)) + \delta\eta(e,(1,1)) \geq (1-\delta)\underline{u}_1 + \delta\eta(e,(0,1)), \tag{3}$$
$$(1-\delta)u_2(e,(1,1)) - \delta\eta(e,(1,1)) \geq (1-\delta)\underline{u}_2 - \delta\eta(e,(1,0)). \tag{4}$$

Proposition 4. *Let W^* be the CEV set.*

1. $\mathrm{Span}(W^*)$ *is equal to the maximal fixed point of* $\Gamma \equiv \gamma_1 + \gamma_2$*, where*

$$\gamma_1(d) \equiv \max_{e_1,f_1,\eta_1}\Big(\pi_1 u_2(e_1,f_1) - \pi_2 u_1(e_1,f_1) - \frac{\delta}{1-\delta}\eta_1(e_1,f_1)\Big) \tag{5}$$
$$\gamma_2(d) \equiv \max_{e_2,f_2,\eta_2}\Big(\pi_2 u_1(e_2,f_2) - \pi_1 u_2(e_2,f_2) + \frac{\delta}{1-\delta}\big(\eta_2(e_2,f_2) - d\big)\Big) \tag{6}$$

where for each $i \in \{1,2\}$, each maximization is subject to $\eta_i : \mathbb{R}_+ \times \{0,1\} \to [0,d]$ and (e_i, f_i) is incentive compatible and individually rational under η_i;

2. $\mathrm{Level}(W^*) = \max_{\bar{e},\bar{f},\bar{\eta}}\big(u_1(\bar{e},\bar{f}) + u_2(\bar{e},\bar{f})\big)$*, subject to $\bar{\eta} : \mathbb{R}_+ \times \{0,1\} \to [0, \mathrm{Span}(W^*)]$ and $(\bar{e},\bar{f})$ is incentive compatible and individually rational under $\bar{\eta}$;*

3. *The endpoints of W^* are*

$$z^1 = (-1,1)\,\gamma_1\big(\mathrm{Span}(W^*)\big) + \pi\,\mathrm{Level}(W^*), \tag{7}$$
$$z^2 = (1,-1)\,\gamma_2\big(\mathrm{Span}(W^*)\big) + \pi\,\mathrm{Level}(W^*). \tag{8}$$

Proof. Follows from Theorem 5 in Miller and Watson (2013). □

This proposition provides an algorithm for computing the CEV set: first solve part 1 to compute the span, then solve part 2 to compute the level, and then calculate part 3 to find the endpoints. In part 1, $\gamma_1(d)$ finds the disagreement behavior that supports an agreement that maximally punishes the agent. The first two terms in the objective function, $\pi_1 u_2(e_1,f_1) - \pi_2 u_1(e_1,f_1)$, value pushing the stage game payoff in a direction (favoring the principal) that is perpendicular to the direction in which the parties will negotiate. The third term, $\frac{\delta}{1-\delta}\eta_1(e_1,f_1)$, values transferring as much continuation value to the principal as possible. The constraint that $\eta_1 : \mathbb{R}_+ \times \{0,1\} \to [0,d]$ ensures that, whatever is the outcome of the stage game, the transfer of continuation value is feasible if the span of the set of continuation values is d. Similarly, $\gamma_2(d)$ finds the disagreement behavior that supports an agreement that maximally punishes the principal. Together, $\Gamma(d) = \gamma_1(d) + \gamma_2(d)$ is the span of continuation values supportable from the start of the current period if the span available starting in the next period is d. When $d = \Gamma(d)$, the span of continuation values is "self-generating" (Abreu, Pearce, and Stacchetti 1990) and consistent with IAC and NFD.

Once $\mathrm{Span}(W^*)$ is known, it is straightforward to find $\mathrm{Level}(W^*)$: simply find the stage game behavior that maximizes the sum of the parties' payoffs, subject to the constraint that it must be incentivized by transfers of continuation value that are feasible given $\mathrm{Span}(W^*)$. This is why $\mathrm{Span}(W^*)$ is the largest fixed point of Γ—it provides the widest scope for incentives, allowing joint payoffs to be maximized subject to IAC and NFD, thereby satisfying PEAC. Finally, if the span, the level, and the behavior under disagreement at each extreme of the CEV set are known, the endpoints of the CEV set are identified.

A.3 Constructing a contractual equilibrium

This appendix section shows how to derive the equilibrium described in Section 4 from Proposition 4. Because we assume $\underline{u} = (0, 0)$, we can assume that both parties engage after every history, since zero effort gives the same payoff as disengagement. The first step is to characterize $\gamma_1(d)$, which determines play under disagreement when punishing the agent.

Lemma 1. *$\gamma_1(d) = \pi_1\sqrt{\frac{\delta}{1-\delta}d}$.*

Proof. To support action profile $(e_1, (1, 1))$, incentive compatibility requires that

$$\eta_1(e_1, (1, 1)) - \eta_1(0, (1, 1)) \geq \frac{1-\delta}{\delta}e_1^2. \tag{9}$$

Since the objective function of $\gamma_1(d)$ is decreasing in $\eta_1(e_1, (1, 1))$, for a given e_1 at best we can set

$$\eta_1(e_1, (1, 1)) = \frac{1-\delta}{\delta}e_1^2 \geq \eta_1(0, (1, 1)) = 0 \tag{10}$$

(ignoring the constraint that $\eta_1(e_1, (1, 1)) \leq d$), in which case the objective function simplifies to

$$\gamma_1(d) = \pi_1 e_1. \tag{11}$$

Evidently it is optimal to maximize e_1, subject to the constraint that $\eta_1(1, (1, 1)) \leq d$. This is solved at $e_1 = \sqrt{\frac{\delta}{1-\delta}d}$, yielding

$$\gamma_1(d) = \pi_1\sqrt{\frac{\delta}{1-\delta}d}. \tag{12}$$

□

The next step is to characterize $\gamma_2(d)$, which determines play under disagreement when rewarding the agent.

Lemma 2. $\gamma_2(d) = 0$.

Proof. To support action profile $(e_2, (1,1))$, incentive compatibility requires that

$$\eta_2(e_2, (1,1)) - \eta_2(0, (1,1)) \geq \frac{1-\delta}{\delta} e_2^2. \tag{13}$$

Since the objective function $\gamma_2(d)$ is increasing in $\eta_2(e_2, (1,1))$, at best we can set

$$\eta_2(e_2, (1,1)) = d \geq \eta_2(0, (1,1)) = d - \frac{1-\delta}{\delta} e_2^2 \tag{14}$$

(ignoring the constraint that $\eta_2(0, (1,1)) \geq 0$), in which case the objective function simplifies to

$$\gamma_2(d) = -\pi_1 e_2 - e_2^2. \tag{15}$$

Evidently it is optimal to minimize e_2, which does not violate the ignored constraint. The solution is to set $e_2 = 0$, yielding $\gamma_2(d) = 0$ for all d. □

Next we compute the span of W^*, which equals the largest fixed point of $\Gamma(d) = \gamma_1(d) + \gamma_2(d)$.

Lemma 3. $\text{Span}(W^*) = \pi_1^2 \frac{\delta}{1-\delta}$.

Proof. We have that $\Gamma(d) = \pi_1 \sqrt{\frac{\delta}{1-\delta} d}$. Its maximal fixed point is computed as follows:

$$d = \pi_1 \sqrt{\frac{\delta}{1-\delta} d} \quad \implies \quad d = \pi_1^2 \frac{\delta}{1-\delta}. \tag{16}$$

□

The next step is to compute the level of W^*.

Lemma 4. $\text{Level}(W^*) = \frac{\delta}{1-\delta} \pi_1$.

Proof. To support action profile $(\bar{e}, (1,1))$, incentive compatibility requires that

$$\bar{\eta}(\bar{e}, (1,1)) - \bar{\eta}(0, (1,1)) \geq \frac{1-\delta}{\delta} \bar{e}^2. \tag{17}$$

Since the objective function is invariant to $\bar{\eta}$, it is optimal to maximize $\bar{e}$ subject to

$$\operatorname{Span}(W^*) \geq \bar{\eta}(\bar{e},(1,1)) \geq \bar{\eta}(\bar{e},(1,1)) - \frac{1-\delta}{\delta}\bar{e}^2 \geq \bar{\eta}(0,(1,1)) \geq 0. \tag{18}$$

This is solved at $\bar{e} = \sqrt{\frac{\delta}{1-\delta}\operatorname{Span}(W^*)}$, yielding

$$\operatorname{Level}(W^*) = \sqrt{\frac{\delta}{1-\delta}\operatorname{Span}(W^*)} = \frac{\delta}{1-\delta}\pi_1. \tag{19}$$

□

Finally, we characterize the CEV set.

Lemma 5. *The CEV set is the line segment with endpoints* $z^1 = \left(0, \frac{\delta}{1-\delta}\pi_1\right)$ *and* $z^2 = \left(\frac{\delta}{1-\delta}\pi_1^2, \frac{\delta}{1-\delta}\pi_1\pi_2\right)$.

Proof. We compute the endpoints as follows:

$$z^1 = (-1,1)\pi_1^2\frac{\delta}{1-\delta} + \pi\frac{\delta}{1-\delta}\pi_1 = \left(0, \frac{\delta}{1-\delta}\pi_1\right), \tag{20}$$

$$z^2 = \pi\frac{\delta}{1-\delta}\pi_1 = \left(\frac{\delta}{1-\delta}\pi_1^2, \frac{\delta}{1-\delta}\pi_1\pi_2\right). \tag{21}$$

□