Relational Contracting, Negotiation, and External Enforcement

David Miller, Trond E. Olsen, and Joel Watson*

March 2018

Abstract

We study relational contracting and renegotiation in environments with external enforcement of long-term contractual arrangements. An external, long-term contract governs the stage games the contracting parties will play in the future (depending on verifiable stage-game outcomes) until they renegotiate. In a *contractual equilibrium*, the parties choose their individual actions rationally, they jointly optimize when selecting a contract, and they take advantage of their relative bargaining power. Our main result is that in a wide variety of settings, in each period of a contractual equilibrium the parties agree to a *semi-stationary* external contract, with stationary terms for all future periods but special terms for the current period. In each period the parties renegotiate to this same external contract, effectively adjusting the terms only for the current period. For example, in a simple principal-agent model with a choice of costly monitoring technology, the optimal contract specifies mild monitoring for the current period but intense monitoring for future periods. Because the parties renegotiate in each new period, intense monitoring arises only off the equilibrium path after a failed renegotiation.

*Miller: University of Michigan; Olsen: NHH Norwegian School of Economics; Watson: UC San Diego. For helpful comments and encouragement, the authors thank Daniel Barron, Andy Skrzypacz, Nageeb Ali; organizers and participants at the 2016 Econometric Society Winter Meetings, the 2016 Stanford Institute for Theoretical Economics summer conference, the Northwestern Third Annual Conference on Relational Contracts (2017), and the 2018 American Economic Association Winter Meetings; as well as seminar participants at Cambridge INET Institute, Florida International University, Indiana University, Johns Hopkins Cary Business School, Kellogg, Konstanz, Michigan, NHH Norwegian School of Economics, Penn State, Princeton, Seoul National University, Shanghai Jiaotong University, Tulane, UCSD, and Washington. Watson thanks the NSF (SES-1227527) and UCSD for grant support.

Long-term contractual relationships are typically governed by a combination of self-enforced arrangements (the parties' coordinated behavior to reward and punish each other over time) and recourse to some degree of external enforcement such as provided by a court system. Macaulay (1963) famously observed that contractual relationships between U.S. firms were often structured with loosely specified legal terms that persisted over time, suggesting the importance of self-enforcement as well the expectation that parties would work things out should disagreements arise (Malcomson 2013). While the literature on *relational contracting* has generated insights on the self-enforced aspects of ongoing contractual relationships, it is important to also investigate the roles of external enforcement and the ever-present opportunity for parties to renegotiate all contractual terms.

This paper presents the first model of ongoing relationships that explicitly accounts for recurring negotiations, self-enforcement, and external enforcement of long-term contractual provisions.[1] We provide a general framework and foundational results for a wide range of settings with moral hazard. The modeling exercise identifies key features of optimal contracting and explains some actual practices, such as the interplay of long-run and short-run contractual provisions, stationary contract terms, and the allocation of control rights.

We view the contract between parties as having two components. The externally enforced part, which we call the *external contract*, prescribes how a court or other external referee is to intervene in the relationship. The self-enforced part, which we call the *regime*, specifies the parties' individual productive actions over time, as well as their anticipated revisions of the external contract.[2] Both the external contract and the regime are renegotiable. Though the productive technology is stationary, the parties' ability to write an arbitrary long-term external contract introduces endogenous non-stationarity: In the current period, the external contract terms agreed upon previously can be changed only by mutual agreement and thus constitute a payoff-relevant state variable. A key question is whether the external contract should specify only stationary terms or should be non-stationary.

The prior literature establishes that, without external enforcement, if the parties can pay monetary transfers that enter their payoffs linearly, then optimal behavior on the equilibrium path is stationary (see, e.g., Levin 2003; Miller and Watson 2013).[3] Introducing external

[1] A recent paper along the same lines, Kostadinov (2017), is discussed below.

[2] In the literature, external and self-enforced (internal) contractual elements are variously differentiated with words such as "explicit/implicit," "formal/informal," and "legal/relational." The terminology we prefer focuses attention on the source of the enforcement power. While the "legal/relational" terminology does so as well, we prefer to think of a "relational contract" as encompassing both an external contract and regime.

[3] In a stationary environment without external enforcement, a relational contract is defined as either a perfect public equilibrium (e.g., Levin 2003) or a contractual equilibrium (Miller and Watson 2013) of an infinitely repeated game. The latter concept explicitly incorporates bargaining and a theory of disagreement.

enforcement, we find that while it is optimal for the contracting parties to write the same external contract every time they renegotiate, the external contract they select is itself non-stationary. If the external enforcer can compel monetary transfers as a function of verifiable outcomes (or if no outcomes are verifiable), then the non-stationarity takes a particular form. Optimally, the long-term part of the external contract, which governs future periods, is stationary; but the short-term part, which governs the current period, is special. We call such a contract *semi-stationary*. Intuitively, the parties choose the long-term part to maximize the power of incentives, while they choose the short-term part to maximize their joint payoffs given the power of incentives available to them. Since they anticipate renegotiating to the same external contract in each new period, along the equilibrium path they always operate under the short-term part of the external contract. Critically, anticipated renegotiation in future periods turns out not to affect the power of incentives.

Allowing for arbitrary long-term externally enforced contracts sets our model apart from the previous literature on relational contracting with limited external enforcement (e.g. Baker, Gibbons, and Murphy 1994, 2002, Schmidt and Schnitzer 1995, Che and Yoo 2001, Kvaløy and Olsen 2009, Iossa and Spagnolo 2011, and Itoh and Morita 2015), which has typically either allowed for only short-term (spot) external enforcement, or assumed that long-term externally enforced contracts are stationary. Moreover, this literature has mostly assumed that self-enforced relational arrangements are irrevocably terminated after a deviation, so then parties behave myopically. In contrast, we suppose that the parties can renegotiate and re-evaluate all aspects of their relationship every period, and we find that they choose to continue with both relational self-enforcement and external enforcement after any history. Our approach thus addresses the question of how agents initiate and manage their relationship, including how their agreements evolve after deviations and disagreements.

The most closely related modeling exercise is Kostadinov (2017), which allows for non-stationary long-term externally enforced contracts and renegotiation. Kostadinov's model builds on Pearce and Stacchetti (1998) and is complementary to ours in that the negotiation theory and equilibrium concepts are different (Kostadinov examines subgame perfect equilibrium without a theory of bargaining power). It is restricted to a simple principal-agent stage game and assumes risk aversion on the part of the agent. Consistent with our modeling exercise, Kostadinov finds that an optimal externally enforced long term contract is renegotiated in equilibrium, and this occurs for reasons similar to those found in our framework.[4] Relations to the literature are discussed further in Section 5.

[4]The continuation contract is designed to allow for harsh punishments to a deviating player, but the contract is then renegotiated ex post to better support efficiency on the equilibrium path.

Our solution concept is *contractual equilibrium* (Miller and Watson 2013), applied to a hybrid repeated game in which each period contains two phases: a cooperative *negotiation phase* and a non-cooperative *action phase*. In the negotiation phase, players can make monetary transfers, and the solution concept predicts they will reach an agreement that satisfies the generalized Nash (1950) bargaining solution. The bargaining set contains all valid continuation payoff vectors; the disagreement point entails no immediate transfer and is determined in equilibrium. In the action phase, the players' actions depend only on the public history and must satisfy individual incentive constraints, just as in a perfect public equilibrium. Since Miller and Watson provide fully non-cooperative foundations using cheap-talk bargaining and axiomatic equilibrium selection, in this paper we restrict attention to the hybrid cooperative/non-cooperative game.[5]

Our modeling approach allows for a broad range of external enforcement capabilities. The external enforcer can impose a stage game for the contracting parties to play, and selection of the stage game can depend on the verifiable outcomes in prior periods. Thus the enforcer's capabilities are defined by the set of stage games it has available to impose, where each stage game includes a partition defining the extent to which the enforcer can verify outcomes.

To illustrate the components of our theory and the main conclusion for contract design, we present in the next section a simple application: a principal-agent relationship with the choice of a costly and externally enforceable monitoring technology. We show that the optimal semi-stationary contract specifies mild monitoring for the current period but intense monitoring for future periods. Since the parties renegotiate in each new period, intense monitoring is enforced only out of equilibrium after a failed renegotiation. Specification of intense monitoring affects disagreement payoffs in such a way that the span of available continuation payoffs, accounting for renegotiation, is enlarged. The larger span enables the parties to save on costly monitoring in the current period.

Following the monitoring application, we present the general model in Section 2 and the analysis of existence, optimal contracts, and semi-stationarity in Section 3 (with technical foundations in the Appendices). Section 4 returns to applications, including an expansion of the monitoring example, a model of multitasking, and a partnership example.

A common theme in the applications is that, because the equilibrium external contracts are semi-stationary, strict contractual terms are routinely renegotiated to milder terms. This implies that the strict terms are actually never imposed in equilibrium. It is noteworthy that

[5] Generalizing Miller and Watson's fully non-cooperative framework to allow for external enforcement would be notationally cumbersome but conceptually straightforward. We comment on this in Section 2.

this type of behavior is often observed in reality. For instance, it is common practice in many organizations to have strict formal rules for employees (e.g., with respect to attendance and procedures at work) but to allow and accept considerable flexibility regarding adherence to these rules. Our framework provides an explanation for such practices.[6]

It is well known that strategic flexibility can be valuable when some, but not all, actions for the players can be externally enforced. We show that such flexibility can be achieved by letting the externally enforced terms of the contract take the form of options. In the monitoring example, allowing the principal to select between strict and mild monitoring improves equilibrium welfare relative to specifying a contractually fixed level of monitoring. Further, decision rights are shown to matter in such settings, and rights tend to be optimally allocated to the party with the highest bargaining power.

Our applications also show that, while the long-term external contract is in general modified through renegotiation each period, this need not be the case in all environments. In a multitask setting, where an agent supplies efforts on two tasks with, respectively, verifiable and non-verifiable but observable outputs, under some conditions the optimal contract utilizes external enforcement for the former task (via a payment schedule) and self-enforcement for the latter task. Further, the externally enforced payment schedule is never renegotiated, but the parties realize that the quantities and payments selected from this schedule will depend on whether they fail to reach agreement.

1 Example: Choice of a Monitoring Technology

For an illustration of the model, consider a relationship between a worker and a manager, where the extent to which the manager can monitor the worker's effort is determined by a costly monitoring technology that can be externally enforced—for instance by a third party who is hired to observe the worker.

The worker (player 1) and the manager (player 2) interact over discrete time periods with an infinite horizon and a shared discount factor δ. Each period comprises two phases:

- the *negotiation phase*, where the players can establish or revise their contract, as well as make immediate, balanced monetary transfers; and
- the *action phase*, where productive interaction occurs.

In the negotiation phase, the immediate net monetary transfer paid from the manager to the worker is denoted $m_1 \in \mathbb{R}$.

[6] Iossa and Spagnolo (2011) provide a related explanation. We discuss the differences in Section 5.

In the action phase, the worker chooses her action a: either low effort ($a = 0$), or high effort ($a = 1$). High effort imposes a personal cost of $\beta \in (0, 1)$ on the worker and yields a benefit of 1 to the manager, both in monetary terms. The players jointly observe a signal x generated by a monitoring technology with accuracy parameter $\mu \in [0, 1]$. If the worker exerts high effort then the signal realization is high ($x = 1$) for sure, but if the worker exerts low effort then the signal realization is either high, with probability $1 - \mu$, or low ($x = 0$), with probability μ. The monitoring technology imposes a cost of $k(\mu)$ on the manager that is strictly increasing in μ and satisfies $\beta + k(1) \leq 1$, so high effort with maximal monitoring generates higher welfare than low effort with minimal monitoring.

At the end of each period, the players publicly observe the signal generated by the monitoring technology. However, only the worker observes her own effort choice a. To keep things simple we also assume that the manager does not observe the payoff he receives in the stage game.[7] We assume that the players can take advantage of arbitrary public randomization devices to coordinate their play. Also, we adopt the standard normalization and multiply the payoffs by $1 - \delta$, which simplifies some expressions and figures.

1.1 Contractual equilibrium with fixed monitoring technology

Suppose first that the monitoring technology μ is fixed exogenously. Before characterizing the contractual equilibrium, let us briefly consider an optimal perfect public equilibrium as analyzed by Levin (2003); in this case, there is no negotiation but players can still make voluntary transfers in the negotiation phase. High effort from the worker and payments from the manager can then be sustained in equilibrium if the cost saved by a deviation is no larger than the expected loss of future surplus, weighted by the probability of detecting the deviation—that is, if $(1 - \delta)\beta \leq \delta\mu(1 - \beta)$. Monitoring costs do not appear because they are fixed irrespective of behavior. This equilibrium can be sustained by reversion to low effort and no payments in all future periods if any party should deviate. However, such behavior is not credible if the parties can renegotiate and can each exercise bargaining power. Contractual equilibrium explicitly accounts for such negotiations.

Since the monitoring technology μ is fixed exogenously, in the negotiation phase the players have only their immediate transfer and their self-enforced continuation play to discuss. If they disagree, then there is no immediate transfer and they coordinate on some continuation play from the action phase, anticipating that they will agree in subsequent

[7] It would be enough to assume that the manager learns his stage-game payoff only after a long delay. Alternatively, one could assume that the manager's payoff depends only on the monitoring signal, equaling 1 if $x = 1$ and $-\frac{\mu}{1-\mu}$ if $x = 0$.

periods. Under disagreement, the worker's effort may be high or low, depending on the history. In the predicted agreement, in contrast, the players coordinate on behavior to maximize the sum of their payoffs subject to the equilibrium constraints, and they make an immediate transfer to divide the surplus relative to disagreement. The parties are endowed with fixed bargaining weights $\pi_1 \geq 0$ and $\pi_2 \geq 0$, satisfying $\pi_1 + \pi_2 = 1$, which determine how the surplus is divided.

Since the environment is stationary, it follows that the players always earn the same sum of continuation payoffs under agreement; let L denote this "joint value." Therefore the set of agreement payoff vectors they can obtain—which we denote by V and call the *value set*—is a line segment of slope -1. Moreover V contains its endpoints. Each endpoint is the payoff vector that arises from a bargaining problem whose disagreement point is achieved by incentive-compatible play in the current period followed by a continuation value selected from V as a function of the realized signal and the outcome of the public randomization device in the current period.

Let z^1 and z^2 be the endpoints of V, where z^1 is the worst continuation value for player 1 and z^2 is the worst for player 2. We determine these endpoints using a recursive formulation, where we fix the line segment from z^1 to z^2 as the feasible continuation values from the next period and then we calculate the extremal continuation values $z^{1\prime}$ and $z^{2\prime}$ that can be supported from the start of the current period. The environment being stationary, we know that $z^{1\prime} = z^1$ and $z^{2\prime} = z^2$ for contractual equilibria (i.e., V must be self-generating).

The disagreement point that achieves the extremal value $z^{1\prime}$ is characterized as follows and displayed in Figure 1: With no transfer, the players coordinate on $a = 1$ being played in the current period. Then, if the signal realization is high, the players coordinate on behavior to achieve continuation value $z^1 + (\rho, -\rho)$. If the signal realization is low then the players coordinate on z^1 from the next period. The value of ρ must be large enough to ensure that the worker does not want to deviate to low effort, knowing that if she does deviate then with probability μ her deviation will be detected and she will be punished:

$$-(1-\delta)\beta + \delta(z_1^1 + \rho) \;\geq\; (1-\delta)\cdot 0 + \mu\delta z_1^1 + (1-\mu)\delta(z_1^1 + \rho).$$

Her incentive constraint simplifies to $\mu\delta\rho \geq \beta(1-\delta)$. It is optimal to pick the smallest possible value of ρ because player 1's expected payoff is increasing in ρ. So we set $\rho = \frac{1-\delta}{\delta} \cdot \frac{\beta}{\mu}$, and the disagreement value (from the current period) is

$$\underline{v}^1 = (1-\delta)(-\beta, 1-k(\mu)) + \delta z^1 + \delta(\rho, -\rho). \tag{1}$$

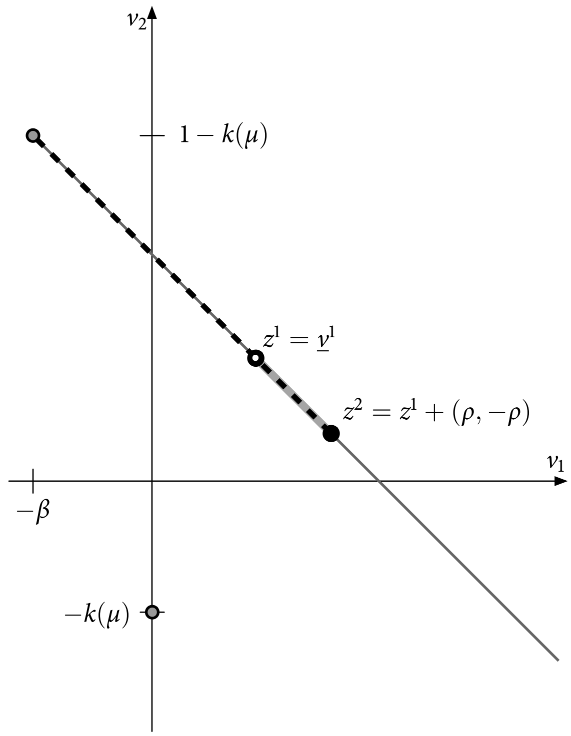

FIGURE 1. CONTRACTUAL EQUILIBRIUM WITH FIXED MONITORING: CONSTRUCTING z^1. All figures in Section 1 are drawn to scale using parameters $\beta = \frac{1}{4}$, $k(\mu) = \frac{1}{2}\mu$, and $\delta = \frac{3}{4}$.

The players can renegotiate from this disagreement point, but it is already efficient so there is no surplus to negotiate over; therefore

$$z^{1\prime} = \underline{v}^1. \tag{2}$$

The disagreement point $z^{2\prime}$ that achieves the extremal value that is worst for player 2 from the current period is characterized as follows, and displayed in Figure 2: With no transfer, the players coordinate on $a = 0$ being played in the current period and, regardless of the signal realization, the players coordinate on behavior to achieve continuation value z^2. Thus, the disagreement value is

$$\underline{v}^2 = (1-\delta)(0, -k(\mu)) + \delta z^2. \tag{3}$$

The players negotiate from this disagreement point to obtain joint continuation value L, and

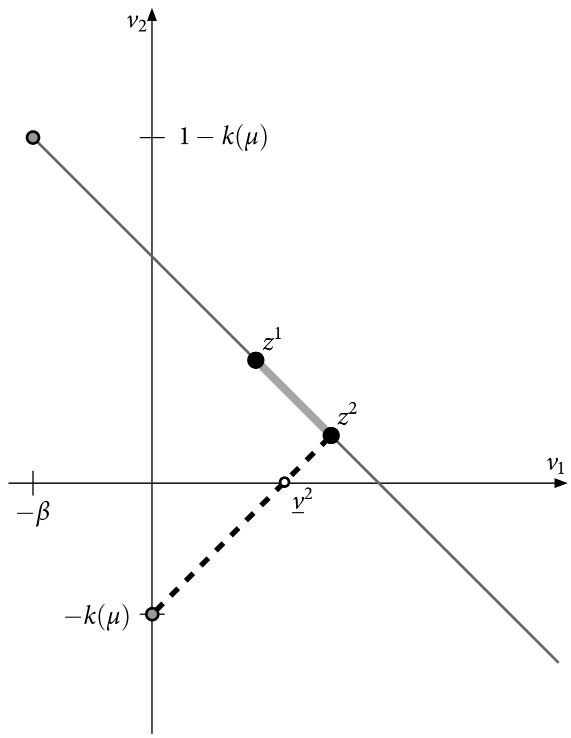

FIGURE 2. CONTRACTUAL EQUILIBRIUM WITH FIXED MONITORING: CONSTRUCTION OF z^2.

they split the surplus according to their bargaining weights π, so we have

$$z^{2\prime} = \underline{v}^2 + \pi(L - \underline{v}_1^2 - \underline{v}_2^2). \tag{4}$$

To complete the calculations, we set $z^{1\prime} = z^1$, $z^{2\prime} = z^2$, and $L = z_1^1 + z_2^1 = z_1^2 + z_2^2$. Making these substitutions and simplifying yields

$$z^1 = \left(\frac{\beta}{\mu} - \beta,\, 1 - k(\mu) - \frac{\beta}{\mu}\right),$$
$$z^2 = (0, -k(\mu)) + \pi(1 - \beta).$$

To interpret these values, note that z_1^1 reflects the worker's rent $(\beta/\mu - \beta)$ when she exerts high effort under imperfect monitoring, plus her share of the surplus relative to disagreeing, which in the case of z^1 is zero because continuation play even under disagreement is efficient. Similarly, z_1^2 reflects the worker's zero rent from exerting zero effort, plus her share of the surplus relative to disagreeing, which in the case of z^2 is $\pi_1(1-\beta)$ because there is zero effort under disagreement.

Note that the *span* of the continuation-value line segment is

$$d \equiv z_1^2 - z_1^1 = z_2^1 - z_2^2 = \pi_1(1-\beta) - (\beta/\mu - \beta)$$

and the level is

$$L = 1 - \beta - k(\mu).$$

This equilibrium outcome requires that $\rho \leq d$; that is, the bonus that the worker receives for a high signal must not exceed the span of V. Recalling that $\rho = \frac{1-\delta}{\delta} \cdot \frac{\beta}{\mu}$, this condition can be expressed in terms of primitives as

$$(1-\delta)\beta \leq \delta\mu\left(\pi_1(1-\beta) - (\beta/\mu - \beta)\right). \tag{5}$$

The condition says that the worker's cost of high effort $(1-\delta)\beta$ can be no greater than the expected future gain, which is $\delta\mu d$. If this inequality does not hold, then high effort cannot be sustained and the contractual-equilibrium value is $(0, -k(\mu))$. Thus if $\mu\pi_1 < 1$ then the condition for sustaining high effort in the contractual equilibrium is stricter than the corresponding condition for the optimal perfect public equilibrium described at the start of this subsection. The difference arises because the perfect public equilibrium employs punishments that are not credible when the parties can renegotiate.

It is important to note how the span and level depend on the monitoring technology μ. The span is increasing in μ, because with better monitoring the worker can be promised a smaller reward ρ for a high signal, which reduces z_1^1. The level is decreasing in μ, because better monitoring costs more. The joint-value maximizing monitoring technology μ^* solves the problem of minimizing $k(\mu)$ subject to Equation 5, which can be written:

$$\mu^* = \frac{1}{\pi_1} \cdot \frac{\beta}{1-\beta}\left(\frac{1-\delta}{\delta} + (1-\mu^*)\right).$$

1.2 Contractual equilibrium with contractible monitoring technology

Now suppose the players can write an external contract that specifies a sequence of monitoring technologies, $\{\mu^t\}$, where μ^t is the level of monitoring to be provided in period t. When the players agree on an external contract, it goes into effect immediately, and it stays in effect until they successfully renegotiate it.

Because the set of feasible contracts is unchanged over time, in each period the parties will make the agreement that attains equilibrium level L^*, regardless of the history of play.

However, the external contract inherited from the most recent prior agreement will still be in force if the parties fail to agree in the current period, and thus it determines what can happen under disagreement. The endpoints of the value set result from agreements formed relative to disagreement play, and therefore depend on the inherited external contract.

It turns out that, in a contractual equilibrium, stationary external contracts (specifying the same μ in all periods) are generally suboptimal. Instead, the optimal external contract is *semi-stationary*, specifying one monitoring level $\hat{\mu}$ for the current period and another level $\tilde{\mu}$ for all future periods. In equilibrium, in each period the inherited contract specifies $\tilde{\mu}$ in all periods, and the parties renegotiate to the same semi-stationary contract with $\hat{\mu}$ for the current period and $\tilde{\mu}$ for all future periods.

Intuition gleaned from the fixed-μ case helps explain this result. To achieve the highest joint value in the current period, the players want μ in this period to be low to save on the monitoring cost. In order to support high effort with a low monitoring level in the current period, the players need the span of continuation values from the next period to be large. To maximize the span, it is best to specify a high monitoring level for future periods, to support wide-ranging disagreement points that will be renegotiated to wide-ranging agreements.

Formally, Equations 1, 3, and 4 are valid for the setting in which the players contract on a sequence of monitoring levels, except that (i) L^* takes the place of L; (ii) z^1 and z^2 depend on the external contract to be inherited in the next period; and (iii) the monitoring level μ in the expressions is what is in force for the current period, not necessarily what will be in force in future periods. In place of Equation 2, we add the following equation, recognizing that the players should renegotiate away from $\underline{v}^1$ if by doing so they can support high effort with monitoring costs lower than specified by the inherited contract:

$$z^{1\prime} = \underline{v}^1 + \pi(L^* - \underline{v}^1_1 - \underline{v}^1_2). \tag{6}$$

Let d be the span of the continuation-value set from any given period $t+1$ and let d' be the span achieved from the start of period t. That is,

$$d \equiv z^2_1 - z^1_1 = z^1_2 - z^2_2 \quad \text{and} \quad d' \equiv z^{2\prime}_1 - z^{1\prime}_1 = z^{1\prime}_2 - z^{2\prime}_2. \tag{7}$$

Putting Equations 1, 3, 4, 6, and 7 together and simplifying yields

$$d' = (1-\delta)\left[\pi_1(1-\beta) - \beta \cdot \frac{1-\mu^t}{\mu^t}\right] + \delta d, \tag{8}$$

To maximize the span from the period t, it is clearly optimal to select $\mu^t = 1$. By induction,

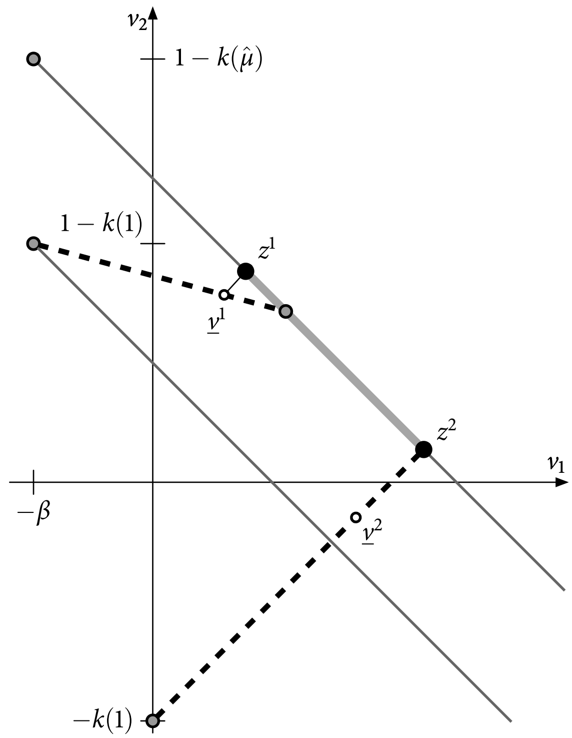

FIGURE 3. CONTRACTUAL EQUILIBRIUM WITH CONTRACTIBLE MONITORING TECHNOLOGY.

to maximize the span by choice of $\{\mu^\tau\}_{\tau=t}^{\infty}$, it is optimal to specify $\mu^t = \mu^{t+1} = \cdots = 1$.

Of course, when they negotiate in period t, the players will want to maximize the span not from period t but from period $t+1$. Therefore they should agree on an external contract that sets $\mu^{t+1} = \mu^{t+2} = \cdots = 1$. This means the span from period $t+1$ solves $d = d'$, which yields $d = \pi_1(1-\beta)$. As for the current period t, to save on monitoring costs that they will actually have to pay today, the parties optimally select the lowest monitoring level that can enforce the worker's high effort. This is the monitoring level for which the worker's required bonus for a high signal, $\rho = \frac{(1-\delta)}{\delta} \cdot \frac{\beta}{\mu}$, just equals the span d. The best choice for μ^t is the smallest value that satisfies this constraint, which is

$$\hat{\mu} = \frac{1}{\pi_1} \cdot \frac{1-\delta}{\delta} \cdot \frac{\beta}{1-\beta}.$$

To summarize, in the contractual equilibrium the players initially choose external contract $\{\mu^t\}$ with $\mu^1 = \hat{\mu}$ and $\mu^t = 1$ for $t = 2, 3, \ldots$. In each subsequent period t, the players revise their contract by specifying $\mu^t = \hat{\mu}$ but leave the specified monitoring level at 1 for all future periods. Note that $\hat{\mu} < \mu^*$ so the players get a strictly higher joint value from a semi-stationary contract than from a stationary contract that has the same μ^* in every period. The contractual equilibrium values are displayed in Figure 3.

2 The Model

We work with a hybrid model as described by Miller and Watson (2013) and Watson (2013), with the addition of an external enforcement technology. Two players $i = 1, 2$ interact in discrete time periods over an infinite horizon with discount factor $\delta \in (0, 1)$. In each period, there are two phases: the cooperative *negotiation phase* and the noncooperative *action phase*. In the negotiation phase, the players make a joint decision to form or revise their contract and make immediate monetary transfers. In the action phase, the players select individual actions and receive payoffs in a stage game. At the end of each period there is also a draw from a public randomization device that we assume is uniformly distributed on the unit interval. We normalize stage-game payoffs by multiplying by $1 - \delta$.

The stage game, which may vary from period to period, is compelled by the external enforcer as directed by the players' *external contract*. The first subsection below formally describes the basic components of the game, including the external enforcement technology. The second subsection details how these components define a relational contracting game. The third subsection describes how we specify a generalized strategy profile for the game, called a *regime*, which include the joint decisions and individual actions. The fourth subsection defines the contractual equilibrium solution concept, which combines individual rationality (self-enforcement) and a theory of bargaining over both the external contract and the self-enforced part of the relationship. Bargaining is resolved according to the generalized Nash bargaining solution, with fixed bargaining weights that represent in reduced form the exogenous parameters of a noncooperative bargaining protocol.

2.1 Technology and external enforcement

Let us describe first the technological details of the relationship, including the scope for external enforcement. A stage game has the following components:

- a set of action profiles $A = A_1 \times A_2$,
- an outcome set X,
- a conditional distribution function $\lambda\colon A \to \Delta X$,
- a payoff function $u\colon A \times X \to \mathbb{R}^2$, and
- a partition P of X.

In the third item, ΔX denotes the set of probability distributions over X. We write $a_i \in A_i$ as player i's individual action in the stage game. The function λ gives the distribution over X for a given action profile. That is, $\lambda(a)$ is the distribution of outcomes in the event that the players select $a \in A$. The outcome $x \in X$ is commonly observed by the players, so each player i knows x and his choice a_i. Player i observes nothing else about actions in the stage game.[8] The partition P represents the external enforcer's verifiability constraints with respect to the stage-game outcome, so that the enforcer can verify only the partition element $P(x)$ containing the realized outcome x.

External enforcement of long term contracts is represented by the following fundamental elements:

- a set G of feasible stage games,
- an abstract set of *external contracts* C,
- a function $g\colon C \to G$,
- an initial external contract $c^0 \in C$, and
- a transition function $\zeta\colon C \times \bigcup_{c \in C} X(c) \times [0,1] \to C$.

These elements describe external enforcement in a convenient recursive formulation. In a given period, an external contract $c \in C$ will be in effect, and $g(c) = (A, X, \lambda, u, P)$ gives the prescribed stage game that the external enforcer compels the parties to play in this period. To make the dependence on c clear, we sometimes write $A(c)$, $X(c)$, $\lambda(\cdot\,; c)$, $u(\cdot\,; c)$, and $P(\cdot\,; c)$ as the components of stage game $g(c)$.

The transition function ζ determines the external contract to be in effect at the beginning of the next period as a function of the current period's external contract, the outcome of the stage game in the current period, and the realization of the public randomization device in the current period. That is, if in the current period the external contract is c, the outcome of the stage game is $x \in X(c)$, and the random draw is $\phi \in [0,1]$, then $\hat{c} = \zeta(c, x, \phi)$ is the external contract in effect at the beginning of the next period. We call $\hat{c}$ the *inherited external contract* for the next period.

To represent the external enforcer's verification constraints, each function $\zeta(c, \cdot\,, \phi)$ must be measurable with respect to the partition $P(\cdot\,; c)$. This means that, for an external contract c, random draw $\phi \in [0,1]$, and any two outcomes $x, x' \in X(c)$ that are in the same partition element (i.e., $x' \in P(x; c)$), we have $\zeta(c, x, \phi) = \zeta(c, x', \phi)$.

[8] In some applications, player i's payoff $u_i(a, x)$ is a function of only x and player i's action $a_i \in A_i$, so that player i obtains no additional information about the other player's actions through her realized payoff. For other applications, we will assume that while $u_i(a, x)$ may depend on the other players' actions, player i does not observe his own payoff.

To get a feel for the formulation, consider as an example an external enforcement technology that allows for arbitrary transitions between stage games as a function of the verifiable outcome. This means that the external contract space is equivalent to the space of functions that map histories to the set of stage games and are measurable with respect to the enforcer's information partitions. A history from period 1 to any given period T is a sequence $\psi = \{\gamma^t, x^t, \phi^t\}_{t=1}^T$, where $\gamma^t = (A^t, X^t, \lambda^t, u^t, P^t) \in G$ denotes the stage game compelled in period t, $x^t \in X^t$ is the outcome, and $\phi^t \in [0,1]$ is the draw of the public randomization device. Let Ψ be the set of all such feasible finite histories, where the case of $T = 0$ is included to denote the null history at the beginning of period 1. Then in this example, an external contract may be defined as any mapping c from Ψ to G that is measurable with respect to the enforcer's information partitions.[9]

2.2 The relational contracting game

We can now describe the contracting game. In each period $t = 1, 2, \ldots$, there are two phases, the first of which is the negotiation phase. Players enter the negotiation phase with an external contract $\hat{c}^t$ that is inherited from the previous period. In the case of $t = 1$, we assume $\hat{c}^1 = c^0$. The players' relative bargaining power in negotiations is fixed for all periods, and is represented by $\pi = (\pi_1, \pi_2)$, where $0 \leq \pi_1 \leq 1$ and $\pi_2 = 1 - \pi_2$. The players negotiate to select an external contract $c^t \in C$ and an immediate monetary transfer $m^t \in \mathbb{R}^2_0$, where $\mathbb{R}^2_0 \equiv \{m \in \mathbb{R}^2 \mid m_1 + m_2 = 0\}$ is the set of real vectors whose components sum to zero (balanced transfers). The negotiated transfer is enforced automatically with the agreement. If the players do not reach an agreement, then $c^t = \hat{c}^t$ and the transfer is zero.[10]

The action phase succeeds the negotiation phase. In the action phase of period t, the players simultaneously choose individual actions in stage game $g(c^t)$, outcome $x^t \in X(c^t)$ is realized according to conditional distribution $\lambda(\cdot\,; c^t)$, and the draw ϕ^t of the public randomization device is realized. Then the external contract inherited in period $t+1$ is $\hat{c}^{t+1} = \zeta(c^t, x^t, \phi^t)$.

The payoffs within a period are given by the sum of any monetary transfer and the stage-

[9]For any stage game $\gamma = (A, X, \lambda, u, P)$, outcome $x \in X$, and public draw ϕ, and for any T-period verifiable history ψ, let $\psi^\frown(\gamma, x, \phi)$ denote the sequence formed by concatenating ψ and (γ, x, ϕ). The transition function ζ is defined so that $\zeta(c, x, \phi) = c(\psi^\frown(g(c), x, \phi))$ for all $\psi \in \Psi$.

[10]Noncooperative foundations can be provided along the lines of Miller and Watson (2013) and Watson (2013). Note that in Miller and Watson there is no external enforcement and so transfers are voluntary. Here we suppose that the immediate transfer is externally enforced, which most easily allows us to extend the noncooperative foundations to the present context. We could alternatively model the immediate transfers as individual actions, but it would complicate the noncooperative foundations.

game payoffs, normalized by $1-\delta$. That is, if the players transfer $m \in \mathbb{R}^2_0$, play action profile a in stage game (A, X, λ, u, P), and get outcome $x \in X$, then the payoff vector for this period is $(1-\delta)(m + u(a, x))$. As the game progresses, the players' behavior (joint actions and individual actions), along with the outcomes of the exogenous random variables, induces a sequence $\{m^t, u^t, a^t, x^t\}_{t=1}^{\infty}$. The realized continuation payoff vector from any period τ is then

$$\sum_{t=\tau}^{\infty} \delta^{t-\tau}(1-\delta)\left(m^t + u^t(a^t, x^t)\right). \tag{9}$$

Because the realized sequence $\{m^t, u^t, a^t, x^t\}$ will be random, the continuation payoff vector is given by the expectation of Equation 9, conditioned on the history prior to time τ and the players' equilibrium.

2.3 Regimes and continuation values

We introduce a generalized notion of strategy, which we call a *regime*, to represent the specification of both individual actions in the action phase and joint decisions in the negotiation phase, after every history. To define a regime, we first must establish notation for histories.

A shared T-period history for the players is a sequence $h = \{(c^t, m^t, x^t, \phi^t)\}_{t=1}^{T}$ with the property that $x^t \in X(c^t)$ for each $t \in \{1, 2, \ldots, T\}$. Here c^t is the external contract and m^t is the transfer jointly chosen by the players in period t. For $t > 1$, if c^t does not equal the inherited external contract $\zeta(c^{t-1}, x^{t-1}, \phi^{t-1})$, then it means that the players renegotiated in period t to change their external contract.[11] The stage-game outcome x^t and the randomization device realization ϕ^t are commonly observed by the players and thus included in the history. The action profile a^t is not included because the players do not commonly observe each others' actions.[12] Let H be the set of all finite histories, including the initial (null) history h^0. Also, for any T-period history $h \in H$, we denote by $\hat{c}(h) = \zeta(c^T, x^T, \phi^T)$ the external contract inherited in period $T+1$ following history h.

A regime $r = (r^{\mathrm{c}}, r^{\mathrm{m}}, r^{\mathrm{a}})$ prescribes joint decision and individual actions as a function of the history. The function $r^{\mathrm{c}} \colon H \to C$ specifies the external contract the players should agree on at the beginning of each period, as a function of the history. The function $r^{\mathrm{m}} \colon H \to \mathbb{R}^2_0$ specifies the associated immediate transfer that the players should agree to.

[11] This accounting of histories does not differentiate between disagreement and agreeing to keep the contractual arrangements unchanged and to make no transfer. Both would be represented by $c^t = \zeta(c^{t-1}, x^{t-1}, \phi^{t-1})$ and $m^t = 0$. The analysis is not affected by whether this distinction is made, and it is simpler to go without it.

[12] In contractual equilibrium, like in perfect public equilibrium, the joint and individual actions on the equilibrium path are measurable with respect to the commonly observed outcome variables.

That is, $(r^{\mathrm{c}}(h), r^{\mathrm{m}}(h))$ prescribes the joint decision in the negotiation phase of the period following history h. Finally, the function $r^{\mathrm{a}}\colon H \times C \times \mathbb{R}_0^2 \to \cup_{c \in C} \Delta A(c)$ gives the mixed action profile as a function of the history to the action phase in any period. That is, if following history h the players jointly choose (c, m) in the current period, then the prescribed action profile for the current period is $r^{\mathrm{a}}(h, c, m) \in \Delta A(c)$. Because we assume that the players randomize independently, $\Delta A(c)$ is taken to mean the uncorrelated distributions over $A(c)$.

The incentive conditions described in the next subsection will be applied to a subset of histories that relate to the regime being evaluated. Specifically, for any regime r, we will look at a set $H(r)$ of histories in which, in each period, either the players made the agreement specified by the regime or there was disagreement. In other words, we are leaving out histories in which the players jointly deviated in the negotiation phase by selecting an external contract or transfer that was not specified by their regime.[13] For any T-period history h and any integer $t \leq T$, let h^t denote the sub-history given by the first t periods of h. Then a given T-period history $h \in H$ is included in $H(r)$ if and only if, for all $t = 1, 2, \ldots, T$, either $c^t = \zeta(c^{t-1}, x^{t-1}, \phi^{t-1})$ and $m^t = 0$, as in disagreement, or $c^t = r^{\mathrm{c}}(h^{t-1})$ and $m^t = r^{\mathrm{m}}(h^{t-1})$. We categorize histories to the action phase of a period similarly. Let $H^{cm}(r)$ be the set of triples (h, c, m) with the property that $h \in H(r)$ and either $(c, m) = (\hat{c}(h), 0)$ or $(c, m) = (r^{\mathrm{c}}(h), r^{\mathrm{m}}(h))$.

Next, for a given regime r, we define continuation values following histories in $H(r)$ and $H^{cm}(r)$. For any history $h \in H(r)$, let $v(h; r)$ be the vector of expected continuation values from the beginning of the period following history h, assuming that the players behave as specified by r from this point in the game. Likewise, for any $(h, c, m) \in H^{cm}(r)$, let $\tilde{v}(h, c, m, a; r)$ be the expected continuation value from the action phase following history (h, c, m) when the individual action profile is $a \in A(c)$ in the current period and the players behave as specified by r in all future periods:

$$\tilde{v}(h, c, m, a; r) = E_{x,\phi}\left[(1-\delta)u(a, x; c) + \delta v(h, (c, m, x, \phi); r)\right],$$

where the expectation is taken with respect to $x \sim \lambda(a; c)$ and $\phi \sim U[0, 1]$. Also, for every $h \in H(r)$, we have

$$v(h; r) = (1-\delta)r^{\mathrm{m}}(h) + \tilde{v}(h, r^{\mathrm{c}}(h), r^{\mathrm{m}}(h), r^{\mathrm{a}}(h, r^{\mathrm{c}}(h), r^{\mathrm{m}}(h)); r).$$

This implicitly involves an expectation calculation over a if $r^{\mathrm{a}}(h, r^{\mathrm{c}}(h), r^{\mathrm{m}}(h))$ is mixed.

[13] Restricting our incentive conditions to histories in $H(r)$ is without strategic consequence, and helps us avoid some technical issues related to existence.

Finally, for any history $h \in H(r)$, let $\underline{v}(h;r)$ denote the *disagreement point* for the bargaining phase in the period following history h. This is the continuation value under the assumption that the players fail to reach an agreement in the current period, and thus $c^t = \hat{c}(h)$ and $m = 0$, but play in the action phase of the current period and all future behavior is specified by the regime r. That is,

$$\underline{v}(h;r) = \tilde{v}(h, \hat{c}(h), 0, r^{\mathrm{a}}(h, \hat{c}(h), 0); r).$$

2.4 Contractual equilibrium

Contractual equilibrium combines two conditions. First, we have the standard sequential rationality condition for individual actions. In the action phase of each period, the players best-respond to each others' actions, given their anticipated behavior in the continuation of the game.

Definition 1. *A regime r is called* **incentive compatible in the action phase** *if for all $h \in H(r)$, $c \in \{r^{\mathrm{c}}(h), \hat{c}(h)\}$, $m \in \mathbb{R}^2_0$, and for each player i and action $a'_i \in A_i(c)$,*

$$\tilde{v}_i(h, c, m, r^{\mathrm{a}}(h, c, m); r) \geq \tilde{v}_i(h, c, m, (a'_i, r^{\mathrm{a}}_{-i}(h, c, m)); r). \tag{10}$$

In words, player i cannot gain by deviating from $r^{\mathrm{a}}_i(h, c, m)$ in the action phase following history h and joint decision (c, m) in the current period. Here we restrict attention to histories in $H(r)$ and, within the current period, the jointly selected external contract being either that specified by the regime or the disagreement point.

The second condition is that in each period the players' joint action in the negotiation phase is characterized by the generalized Nash bargaining solution, with fixed bargaining weights given by $\pi = (\pi_1, \pi_2)$. That is, the players reach an agreement that maximizes their joint value and they split the bargaining surplus according to their bargaining weights. Note that π is a parameter of the bargaining solution; it summarizes in reduced form the parameters of a corresponding noncooperative bargaining protocol.

Importantly, we assume that the players negotiate over both the external part of the contract and the self-enforced part. The former amounts to the selection of c and an immediate transfer. The latter means coordinating on a regime for the continuation of the game, which includes individual actions in the current and future periods as well as anticipated joint decisions in future periods (all as a function of the history). Following Miller and Watson (2013), we capture this condition by first imposing an *internal consistency* agreement condition, which represents the following idea: In equilibrium, the players recognize that, after any history $h \in H(r)$, they have the option of agreeing to continue as though the history is any other $h' \in H(r)$. That is, the players have the option of selecting the external contract that they would have selected following h' and then plan to play as their regime specifies from history h'.[14] Since the players can make any transfer in the negotiation phase, they are able to split the negotiation surplus in any way desired, and our bargaining assumption implies that they split the surplus according to π.

Definition 2. *A regime r is said to be* **internally bargain-consistent** *if for all $h \in H(r)$, $v_1(h;r) + v_2(h;r) \geq \underline{v}_1(h;r) + \underline{v}_2(h;r)$ and*

$$v(h;r) = \underline{v}(h;r) + \pi \max_{h' \in H(r)} \left(v_1(h';r) + v_2(h';r) - \underline{v}_1(h;r) - \underline{v}_2(h;r) \right).$$

The following lemma follows directly from the definition of internal bargain-consistency.

Lemma 1. *If regime r is internally bargain-consistent, then it has the same joint value from the beginning of any period. That is, there exists $L \in \mathbb{R}$ such that $v_1(h;r) + v_2(h;r) = L$ for all $h \in H(r)$.*

For a regime that is internally consistent, let us call L its *joint value*, *welfare level*, or just *level*. The players jointly prefer to coordinate on a regime that maximizes L, and this condition completes the definition of contractual equilibrium:

Definition 3. *Given exogenous bargaining weights π, a regime is called a* **contractual equilibrium (CE)** *if it is incentive compatible in the action phase and internally bargain-consistent, and its level is maximal among the set of regimes with these properties.*

The recursive formulation of equilibrium is provided in the Appendices. Existence is addressed in the context of our main characterization results in the next section.[15] At this point, we have the following obvious implication of the contractual-equilibrium definition.

Lemma 2. *For a given relational-contract setting, all contractual equilibria attain the same level.*

We conclude this section by observing that strengthening the external enforcement technology implies a higher welfare level in contractual equilibrium. The external enforcement technology becomes stronger if, for instance, the information partitions in the stage games

[14]Note that this is feasible because, just after selecting $r^c(h')$ and any transfer, the continuation game would be the same as from the action phase following h'.

[15]Additionally, an existence result for settings with a finite number of external contracts is provided in Supplementary Appendix B.1.

become finer (and thus external contracts can become "less incomplete"), if the set of feasible stage games expands, or if the set of external contracts expands. Recalling that the external enforcement technology is given by (G, C, g, c^0, ζ), we can relate two technologies most simply by inclusion: *Technology* $(\tilde{G}, \tilde{C}, \tilde{g}, \tilde{c}^0, \tilde{\zeta})$ *is stronger than is technology* (G, C, g, c^0, ζ) if $G \subset \tilde{G}$, $C \subset \tilde{C}$, $\tilde{c}^0 = c^0$, and, when restricted to contracts in C and outcomes in $X(c)$, we have $\tilde{g} = g$ and $\tilde{\zeta} = \zeta$. In this sense, to get a stronger technology we enlarge the set of stage games and external contracts, so all of the items in the weaker technology are retained.

Theorem 1. *If one external enforcement technology is stronger than another, and if a contractual equilibrium exists under both technologies, then the welfare level is weakly higher under the stronger technology.*

Proof. Suppose regime r is incentive compatible in the action phase and internally bargain-consistent under the weaker technology. For the stronger technology, define regime $\hat{r}$ so that $\hat{r}^{\mathrm{c}}(h) = r^{\mathrm{c}}(h)$, $\hat{r}^{\mathrm{m}}(h) = r^{\mathrm{m}}(h)$, and $\hat{r}^{\mathrm{a}}(h, c, m) = r^{\mathrm{a}}(h, c, m)$ for all $h \in H(r)$, $c \in \{r^{\mathrm{c}}(h), \hat{c}(h)\}$, and $m \in \mathbb{R}^2_0$. For other histories, $\hat{r}$ can be arbitrary. We have $H(\hat{r}) = H(r)$ and, by definition, $\hat{r}$ replicates r and is incentive compatible in the action phase and internally bargain-consistent under the stronger technology. Because contractual equilibrium achieves the maximal level over regimes with these properties, the result is implied.[16] □

This conclusion contrasts with some of the prior literature in relational contracts, which has found that under certain assumptions on equilibrium selection, improving external enforcement can harm the contracting parties. The key assumption behind the prior literature's result is that (as in Baker, Gibbons, and Murphy 1994, 2002 and Schmidt and Schnitzer 1995), after any deviation the parties permanently discontinue self-enforced relational arrangements and, instead, in all future periods they play a stage game equilibrium under an optimal externally enforced spot contract. In contrast, contractual equilibrium posits that the parties can always renegotiate both the external contract and their self-enforced arrangements (the regime). Thus, when they successfully renegotiate following any history, the parties agree to an optimal combination of externally enforced and self-enforced elements. Theorem 1 is in line with recent empirical studies that find complementarity between externally enforced and self-enforced contracts (Beuve and Saussier 2012; Lazzarini, Miller, and Zenger 2004; Ryall and Sampson 2009; Poppo and Zenger 2002).

[16] The notion of stronger technology can be expanded to relate technologies that are not ordered by inclusion, without affecting the result. For instance, suppose $\tilde{G}$ includes a stage game that has a strictly finer partition for verification than does a comparable stage game in G; then it is not necessary for $\tilde{G}$ to contain the latter stage game as well. Details are provided in Supplementary Appendix B.2.

3 Optimal Contracts and Semi-Stationarity

In a contractual equilibrium, in each period the players renegotiate their external contract and regime to maximize their joint value, accounting for the fact that these will be renegotiated again when the next period starts. In principle, the optimal external contract may be complicated, specifying different stage games after different histories in order to punish and reward the players for their past behavior. However, we show in this section that simple contracts are optimal in a broad range of settings.

3.1 Stationary and semi-stationary contracts

Consider two categories of simple external contracts:

Definition 4. *An external contract $c \in C$ is called **stationary** if $\zeta(c, x, \phi) = c$ for all $x \in X(c)$ and $\phi \in [0, 1]$.*

Definition 5. *An external contract $c \in C$ is **semi-stationary** if there is a stationary external contract $\underline{c}$ such that $\zeta(c, x, \phi) = \underline{c}$ for all $x \in X(c)$ and $\phi \in [0, 1]$. In this case, say that c **transitions to** $\underline{c}$.*

A stationary external contract c always transitions back to itself, so it specifies the same stage game $g(c)$ in every period regardless of the history. A semi-stationary external contract c starts with stage game $g(c)$ and specifies stage game $g(\underline{c})$ in all future periods regardless of the history.

Note that the form of stationarity described in these definitions pertains to only the external contract. The regime still may specify behavior that changes over time and is sensitive to the history.

Definition 6. *A regime $r = (r^{\mathrm{c}}, r^{\mathrm{m}}, r^{\mathrm{a}})$ is **semi-stationary** if there is a semi-stationary external contract c such that $r^{\mathrm{c}}(h) = c$ for all $h \in H(r)$.*

In a semi-stationary regime, the players always negotiate to the same semi-stationary external contract c that transitions to stationary external contract $\underline{c}$. Although in such a regime the inherited external contract in each period $t > 1$ is the stationary contract $\underline{c}$ that calls for stage game $g(\underline{c})$ in every period, in each period the players negotiate back to the semi-stationary external contract c and therefore along the equilibrium path they play stage game $g(c)$ in every period.

As we show in the next two subsections, in typical contractual settings it is optimal for the players to select semi-stationary external contracts. To be more precise, there are semi-stationary contractual equilibria. Some technical conditions are required for the result, starting with an assumption that C contains all possible semi-stationary external contracts:

Assumption 1. *For every pair of stage games $\gamma, \gamma' \in G$, there is a semi-stationary external contract $c \in C$ that transitions to some stationary contract $c' \in C$ with the property that $g(c) = \gamma$ and $g(c') = \gamma'$.*

As an example, the enforcement technology that allows for arbitrary selection of the stage game as a function of the verifiable history (the setting described at the end of Section 2.1) satisfies this assumption. We also make the following mild assumptions to help ensure existence.

Assumption 2. *The initial external contract c^0 is stationary and $g(c^0)$ has a Nash equilibrium.*

Assumption 3. *The games in G have uniformly bounded joint values: There is a number $\vartheta \in \mathbb{R}$ such that for every stage game $(A, X, \lambda, u, P) \in G$ and every $a \in A$, we have $-\vartheta \leq u_1(a) + u_2(a) \leq \vartheta$.*

We finally note here that, while the semi-stationary regime entails renegotiation from contract $\underline{c}$ to c every period, the same behavior could be implemented in a stationary manner using messages between the contracting parties and the enforcer. The external contract could then specify $\underline{c}$ as a default option that would be enforced if, say, no messages were sent just after the negotiation phase (indicating disagreement); and c as an option that would be enforced if the parties just after the negotiation phase sent a joint message indicating agreement. While such a message contingent contract would be stationary and not renegotiated in equilibrium, we prefer to work with the contracts without messages, since they better highlight the intertemporal changes in contract terms that are actually implemented in equilibrium.

3.2 Semi-stationary with transfers

In many settings the external enforcer can compel arbitrary transfers as a function of the enforcer's verifiable information about the stage-game outcome.

Definition 7. *The contractual setting has* ***externally enforced transfers*** *if for every stage game $(A, X, \lambda, u, P) \in G$ and every P-measurable function $b : X \to \mathbb{R}^2_0$, it is the case that $(A, X, \lambda, u + b, P) \in G$ as well.*

To understand this definition, think of b as specifying a monetary transfer between the players as a function of the verifiable outcome $x \in X$. Transforming a stage game by changing the payoff function from u to $u + b$ is equivalent to adding an externally enforced transfer, which is feasible because b is P-measurable.

Our main result is that, under some technical conditions sufficient for existence, semi-stationary contracts are optimal in settings with externally enforced transfers. Let us start by developing intuition and the basic conceptual argument. We will use some recursive techniques that relate rational behavior in one period to sets of continuation values from the start of the next period. We will also introduce some additional notation needed to state the result formally.

The main logic

For any external contract c, let $W(c)$ be the set of continuation values that can be supported in a contractual equilibrium from the start of a period with inherited contract c. Assume for now that $W(c)$ is nonempty. Because all contractual equilibrium continuation values attain the same level L, we know that $W(c)$ is a subset of the line $\{w \in \mathbb{R}^2 \mid w_1 + w_2 = L\}$ for every $c \in C$. Let us presume that every set $W(c)$ is bounded and thus has a finite *span*, which is defined as the horizontal (equivalently, vertical) distance between the endpoints of the set. Further, suppose that the span of $W(c)$ is maximized by some external contract $\tilde{c}$ and that $W(\tilde{c})$ contains its endpoints z^1 and z^2, where z^1 denotes the value that is worst for player 1 and z^2 is the worst value for player 2. Thus, the span of $W(\tilde{c})$ is $z_1^2 - z_1^1 = z_2^1 - z_2^2$.

Suppose that, following a given history, the players would optimally select external contract c^t which has them play stage game $g(c^t) = (A, X, \lambda, u, P)$ in the current period t. Incentives in the action phase are influenced both by the stage-game payoffs of $g(c^t)$ and by the continuation value v^{t+1} starting in period $t + 1$. Note that v^{t+1} is contingent on the outcome of the action phase in period t. Thus, we can write $v^{t+1} = z(x, \phi)$ for some function $z\colon X \times [0, 1] \to \mathbb{R}^2$. In the action phase of period t, the players are essentially playing the artificial game that has the space A of action profiles and payoffs given by

$$U(a) \equiv E_{x,\phi}[\,(1 - \delta)u(a, x) + \delta z(x, \phi)\,], \tag{11}$$

for each $a \in A$, where the expectation is taken with respect to $x \sim \lambda(a; c^t)$ and $\phi \sim U[0, 1]$.

Figure 4 illustrates the mapping z for stage-game outcomes x and x' that are in the same partition element $P(x)$, so that the external enforcer cannot distinguish between them. For any realization ϕ of the randomization device, the inherited external contract c in period $t+1$ must be the same for x and x', and so both $z(x, \phi) \in W(c)$ and $z(x', \phi) \in W(c)$. However,

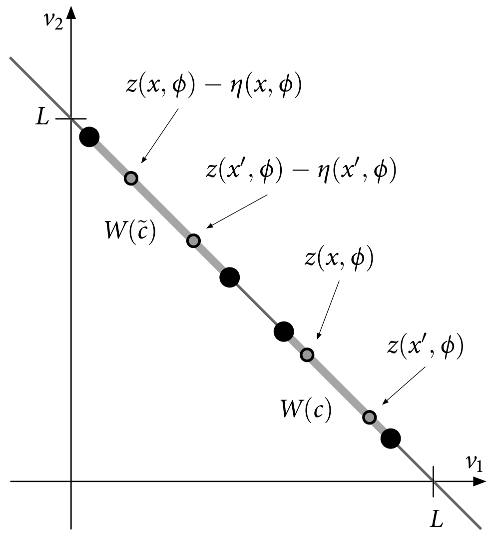

FIGURE 4. THE MAIN LOGIC OF SEMI-STATIONARITY.

the span of $W(c)$ is less than the span of $W(\tilde{c})$, which is located elsewhere on the line of level L.

Let us examine how we can transform u and z to utilize only continuation values in $W(\tilde{c})$, along with externally enforced transfers, *without altering the artificial game* $\langle A, U\rangle$. Clearly we must have $z(x,\phi) \in W(\zeta(c^t, x, \phi))$ for all $x \in X$ and $\phi \in [0,1]$. That is, $z(x,\phi)$ must be a contractual-equilibrium continuation value associated with the external contract inherited in period $t+1$. Because $W(\tilde{c})$ has the greatest span, we can find a function $\eta\colon X \times [0,1] \to \mathbb{R}^2_0$, representing transfers between the players, such that

$$z(x,\phi) - \eta(x,\phi) \in W(\zeta(c^t, x, \phi)) - \eta(x,\phi) \subset \operatorname{Conv} W(\tilde{c})$$

for all $x \in X$ and $\phi \in [0,1]$. Here "Conv" denotes the convex hull. The construction is shown in Figure 4; because the span of $W(\tilde{c})$ is larger than the span of $W(c)$, the former accommodates the distance between $z(x,\phi) - \eta(x,\phi)$ and $z(x',\phi) - \eta(x',\phi)$.

Further, taking the expectation over ϕ, we have

$$E_\phi[z(x,\phi) - \eta(x,\phi)] \in \operatorname{Conv} W(\tilde{c})$$

for all x. We can then find a function $z'\colon X \times [0,1] \to \mathbb{R}^2$ with two key properties: first,

$z'(x, \phi) \in \{z^1, z^2\}$ for all x and ϕ; and, second,

$$E_\phi[z'(x, \phi)] = E_\phi[z(x, \phi) - \eta(x, \phi)]$$

for every x. That is, to achieve the specified expected continuation value in $\mathrm{Conv}\, W(\tilde{c})$ for any particular $x \in X$, we can randomize over the endpoints of $W(\tilde{c})$ using the public random draw ϕ to achieve the needed probabilities.

The foregoing analysis allows us to rewrite Equation 11 by substituting in the new function z', which maps into the set $\{z^1, z^2\}$:

$$U(a) = E_{x,\phi}\big[(1 - \delta)u(a, x) + \delta\eta(x, \phi) + \delta z'(x, \phi)\big]. \tag{12}$$

The final step is to define a function $b\colon X \to \mathbb{R}^2_0$ by setting $b(x) = E_\phi\big[\frac{\delta}{1-\delta}\eta(x, \phi)\big]$ for all $x \in X$. Then we can define payoff function u' by $u'(a, x) = u(a, x) + b(x)$ and, substituting for this in Equation 12, we obtain:

$$U(a) = E_{x,\phi}[\,(1 - \delta)u(a, x) + \delta z(x, \phi)\,] = E_{x,\phi}[\,(1 - \delta)u'(a, x) + \delta z'(x, \phi)\,]. \tag{13}$$

Note that replacing u with u' is possible because, by the assumption of externally enforced transfers, $(A, X, \lambda, u', P) \in G$.

To summarize, we transformed the stage game for period t by adding the transfer function b. Correspondingly, we changed the specification of continuation values in period $t+1$, now given by the function z' that maps to only the endpoints of $W(\tilde{c})$. In other words, whereas the original specification used a selection of various sets of continuation values $W(c)$ in period $t + 1$ to motivate the individual actions in period t, we replaced this selection with transfers at the end of period t. This was possible because (i) continuation-value sets have the same level, so varying between them is equivalent to making a transfer; and (ii) both the externally enforced transfer at the end of period t and the selection of the inherited external contract in period $t + 1$ are conditioned on the verifiable outcome of the stage game in period t. In the end, the artificial game being played in the action phase of period t is unchanged, so the same behavior and continuation values are supported.

In terms of external contract specifications, these adjustments would be accomplished by replacing the given external contract c^t with an external contract $c' \in C$ for which $g(c') = (A, X, \lambda, u', P)$ and $\zeta(c', x, \phi) = \tilde{c}$ for all $x \in X$ and $\phi \in [0, 1]$. Importantly, the transition c' specifies from the current period t to period $t + 1$ is "noncontingent"; that is, it specifies same inherited external contract $\hat{c}^{t+1} = \tilde{c}$ regardless of the outcome of the action phase in period t.

We can repeat the argument with $\tilde{c}$ in period $t+1$. That is, we can find a way to change the stage game and selection of continuation values transition from period $t+1$ to period $t+2$ so that the transition from $\tilde{c}$ to the inherited external contract in period $t+2$ is noncontingent. Proceeding by induction, we can identify an external contract that is noncontingent in every transition, and that achieves the same continuation values as does the original external contract c^t.

Holding aside whether the external contract identified by this procedure is actually an element of C, we can go further by recognizing that a key step in the above logic is finding a set $W(\tilde{c})$ that has the largest span among sets of equilibrium continuation values. Because stage game $g(\tilde{c})$ specified in period $t+1$ is instrumental in achieving the largest span, we would expect that it would be useful to utilize it not just in period $t+1$ but in future periods as well. This is indeed the case, and it implies the optimality of a semi-stationary contract that specifies the same stage game in all periods $t+1, t+2, \ldots$.

The main result

We now describe the necessary technical conditions and the main result. Consider a given period with stage game $\gamma = (A, X, \lambda, u, P) \in G$ and let $y \colon X \to \mathbb{R}^2$ specify the expected continuation value from the next period as a function of the outcome in the current period. The function y will take the place of z that was featured in the preceding analysis. Essentially, y is the expectation of z over the random draw ϕ, and we can now think of y as a choice variable in the optimization problems that will characterize contractual equilibrium.

Let $\overline{u}^\gamma(a) \equiv E_x[\, u(a,x)\,]$ and let $\overline{y}^\gamma(a) \equiv E_x[\, y(x)\,]$, where these expectations are taken with respect to $x \sim \lambda(a)$. We will write A^γ and X^γ as the sets of action profiles and outcomes in stage game γ. Then in the action phase of the current period, the players are essentially playing the artificial game $\langle A^\gamma, (1-\delta)\overline{u}^\gamma + \delta \overline{y}^\gamma \,\rangle$ and incentive compatibility is given by the following definition.

Definition 8. *Given $\gamma \in G$ and $y \colon X^\gamma \to \mathbb{R}^2$, call action profile $\alpha \in \Delta A^\gamma$* ***enforced*** *(relative to γ and y) if it is a Nash equilibrium of $\langle A^\gamma, (1-\delta)\overline{u}^\gamma + \delta \overline{y}^\gamma \,\rangle$.*

We will characterize the span that can be generated for continuation values at the beginning of the current period, for a given stage game γ. Because negotiation will lead to a constant level, we normalize the continuation values from the action phase so that they lie on the line $\mathbb{R}^2_0$ (zero joint value). The normalization is done by shifting stage-game payoffs along the ray π. This translates a payoff vector (u_1, u_2) to $(\pi_2 u_1 - \pi_1 u_2,\ \pi_1 u_2 - \pi_2 u_1)$. We also normalize expected continuation values from the next period to be on the line segment

$$\mathbb{R}^2_0(d) \equiv \{m \in \mathbb{R}^2 \mid m_1 + m_2 = 0 \text{ and } m_1 \in [0, d]\},$$

for a given span d. The justification for considering this full (convex) line segment is that the public draw ϕ can be used to randomize between any values that can be achieved in the next period.

We want to maximize the span of the induced set of continuation values from the current period (written below as the difference between player 1's best and worst continuation values) by choice of the stage game and action profiles. That is, we look for a stage game γ and action profiles α^1 and α^2, where α^1 supports a continuation value that is worst for player 1 and α^2 supports a continuation value that is best for player 1 (worst for player 2). These action profiles must be enforced relative to the stage game and some selection of continuation values from the start of the next period. For any stage game γ, action profile $\alpha \in \Delta A^\gamma$, and function $y\colon X^\gamma \to \mathbb{R}^2_0(d)$, define

$$\omega(\alpha, \gamma, y) = (1 - \delta)\left(\pi_2 \overline{u}^\gamma_1(\alpha) - \pi_1 \overline{u}^\gamma_2(\alpha),\, \pi_1 \overline{u}^\gamma_2(\alpha) - \pi_2 \overline{u}^\gamma_1(\alpha)\right) + \delta \overline{y}^\gamma(\alpha).$$

This is the normalized continuation value. Then let $\Lambda(d)$ denote the maximized difference of player 1's continuation values in a stage game, by choice of the stage game and enforced action profiles:

$$\begin{array}{rrl} \Lambda(d) \equiv & \max & \omega_1(\alpha^2, \gamma, y^2) - \omega_1(\alpha^1, \gamma, y^1) \\ & \text{by choice of:} & \gamma \in G;\ y^1, y^2\colon X^\gamma \to \mathbb{R}^2_0(d);\ \text{and } \alpha^1, \alpha^2 \in \Delta A^\gamma, \\ & \text{subject to:} & \alpha^1 \text{ is enforced relative to } \gamma \text{ and } y^1, \\ & & \alpha^2 \text{ is enforced relative to } \gamma \text{ and } y^2. \end{array} \tag{14}$$

A second optimization problem is to maximize the joint payoff attained in the current period, by choice of the stage game and action profile. For the enforcement condition, we normalize the continuation values from the next period to have joint values of zero.

$$\begin{array}{rrl} \Xi(d) \equiv & \max & \overline{u}^\gamma_1(\alpha) + \overline{u}^\gamma_2(\alpha) \\ & \text{by choice of:} & \gamma \in G,\ y\colon X^\gamma \to \mathbb{R}^2_0(d),\ \text{and } \alpha \in \Delta A^\gamma, \\ & \text{subject to:} & \alpha \text{ is enforced relative to } \gamma \text{ and } y. \end{array} \tag{15}$$

Our main result establishes that, assuming that Optimization Problems 14 and 15 have solutions, a contractual equilibrium exists and is semi-stationary.

Theorem 2. *Suppose Assumptions 1–3 hold and the contractual setting has externally enforced transfers. If $\Lambda(d)$ and $\Xi(d)$ exist for all $d \geq 0$ then there exists a semi-stationary contractual equilibrium. The contractual-equilibrium span d^* is the largest fixed point of Λ, which exists, and the level is $L^* = \Xi(d^*)$.*

Since by definition all contractual equilibria in the game attain the same joint payoffs, this theorem shows that semi-stationarity is optimal. The proof, provided in Appendix A.3 with the technical foundations developed in Appendix A.1 and Appendix A.2, shows how to construct the contractual equilibria by calculating the solutions to Programs 14 and 15.

3.3 Semi-stationarity with no verifiable information

Next consider settings in which the external enforcer cannot distinguish between any stage-game outcomes.

Definition 9. *The contractual setting is said to have* ***no verifiable information*** *if for every $g = (A, X, \lambda, u, P) \in G$, the partition P is trivial: $P = \{X\}$.*

Without verifiable information, an external contract specifies the sequence of stage games to be played but cannot make the sequence conditional on the history of stage-game outcomes. For instance, the example in Section 1 has no verifiable information, because the external enforcer does not observe the monitoring signal. The following result shows that semi-stationarity is optimal when there is no verifiable information, even if the external contracting authority will not compel transfers.

Theorem 3. *Suppose Assumptions 1–3 hold and the contractual setting has no verifiable information. If $\Lambda(d)$ and $\Xi(d)$ exist for all $d \geq 0$ then there exists a semi-stationary contractual equilibrium. The contractual-equilibrium span d^* is the largest fixed point of Λ, which exists, and the level is $L^* = \Xi(d^*)$.*

Proof. We can prove this theorem by transforming the contracting environment into one to which Theorem 2 applies. For any relational contract setting, augment G so that there are externally enforced transfers. This will change neither the contractual-equilibrium set nor Optimization Problems 14 and 15, because externally enforced transfers cannot be conditioned on the outcome of the action phase in any period. In other words, externally enforced transfers can be only constants, which coincide with what the players can do voluntarily in the course of bargaining in each period. From Theorem 2, we have a semi-stationary contractual equilibrium. If such an equilibrium specifies selection of non-zero externally enforced transfers, it is straightforward to replace these transfers with voluntary transfers in the bargaining phase and the equilibrium conditions remain satisfied. □

4 More Examples

In this section we take advantage of the generality of our framework to explore several examples. The first builds on the monitoring example from Section 1 by adding a verifiable action—an option choice—that the monitoring technology and externally enforced transfers can be conditioned on. The ability to write an option contract improves matters for the parties without changing the underlying production technology. Moreover, the optimal assignment of the option (to the worker or manager) depends on relative bargaining power.

The second example generalizes the underlying technology, enabling the worker to exert effort on both quantity and quality dimensions. The external enforcer can verify quantity, but not quality, and enforces transfers that depend on quantity. As in the first example, an option contract is optimal, and who should have decision rights over the option depends on relative bargaining power.

The third example considers a partnership with no verifiable information, but in which the external contract can specify a parameter that changes the production technology, specifically the relative payoffs of cooperation and a single player shirking. This example generalizes the prisoners' dilemma example from Miller and Watson (2013), in which, without external enforcement, the partnership cannot succeed. With external enforcement, even without verifiable information, the parties can mitigate their moral hazard problems by specifying that they should employ a safer but less efficient technology under disagreement.

For each example, a sketch of the analysis is included here. Supplementary Appendices B.3-B.5 contain additional details.

4.1 Option contracts and the allocation of decision rights

Consider again the monitoring example, but suppose now that both the monitoring level μ and an associated transfer $p = (-p, p) \in \mathbb{R}^2_0$ can be externally enforced. Further, let us assume that the action phase includes an option choice for the manager: First the manager chooses from a menu of two monitoring/payment pairs, (μ^1, p^1) and (μ^2, p^2); then the external enforcer compels the chosen monitoring technology μ^j and the transfer p^j from the worker to the manager; and finally the agent selects effort. The external contract specifies the menu items (μ^1, p^1) and (μ^2, p^2) for each period, so the set of stage games is given by the feasible two-option menus, $(p^1, p^2, \mu^1, \mu^2) \in \mathbb{R}^2 \times [0, 1]^2$. The external contract identified in Section 1 can be written here by selecting $\mu^1 = \mu^2$ and $p^1 = p^2$, and it is easy to see that a fixed payment will not affect the equilibrium outcome.

Optimal option contract

We show here that allowing option contracts strictly increases the equilibrium welfare level compared to the level that can be attained with no externally enforced payments. The flexibility to write an option contract allows for the worker to be treated differently under disagreement when she is to be rewarded versus when she is to be punished. The optimal external contract is, of course, semi-stationary, with the same menu items (μ^1, p^1) and (μ^2, p^2) specified for every future period. For the current period, parties select a specific monitoring level $\hat{\mu}$ (technically the two options specify the same $\hat{\mu}$ and no payment) to maximize their attainable joint value.

In the stationary part of the external contract, under disagreement the manager's option choice depends on whether the worker is being punished or rewarded. For a history in which the worker is being punished, the manager will select option (μ^1, p^1), which entails high monitoring and a high payment from the worker. In this case the worker will exert high effort ($a^1 = 1$). In contrast, when the worker is being rewarded the manager will select option (μ^2, p^2) entailing low monitoring and a low payment from the worker, and in this case the worker will exert low effort ($a^2 = 0$). In an equilibrium regime, these choices must be incentive compatible, should disagreement occur. It must also be the case that negotiations at the beginning of each period lead to agreements where gains from agreeing relative to disagreeing are split according to the parties' bargaining weights.

The analysis proceeds as in Section 1. See Supplementary Appendix B.3 for details. Given the behavior just described for the disagreement outcomes that most favor the worker and most favor the manager, it turns out that, incorporating renegotiation, the worker's worst (z_1^1) and best (z_1^2) continuation payoffs are given by

$$z_1^j = -p^j + (\beta/\mu^j - \beta)a^j + \pi_1(L - L^j), j = 1, 2, \tag{16}$$

where $L^j = (1 - \beta)a^j - k(\mu^j)$, $j = 1, 2$ gives the welfare levels under disagreement when the worker is being punished ($j = 1$) and rewarded ($j = 2$); and L is the level under agreement.[17] The payoffs reflect the worker's payment p^j, her rent $\beta/\mu^j - \beta$ when she exerts effort under imperfect monitoring, and her share of the welfare gains from agreeing relative to disagreeing.

The incentive-compatibility condition for the manager to select option (μ^2, p^2) under disagreement when rewarding the worker is $p^2 - k(\mu^2) \geq p^1 - k(\mu^1)$.[18] This implies that

[17] These payoffs coincide with those given in Section 1 when $p^j = 0$ and $\mu^j = \mu$, $j = 1, 2$.

[18] The corresponding condition for the other disagreement point does not bind. The regime specifies that if the manager selects the wrong option then the worker should exert zero effort and the parties will coordinate to

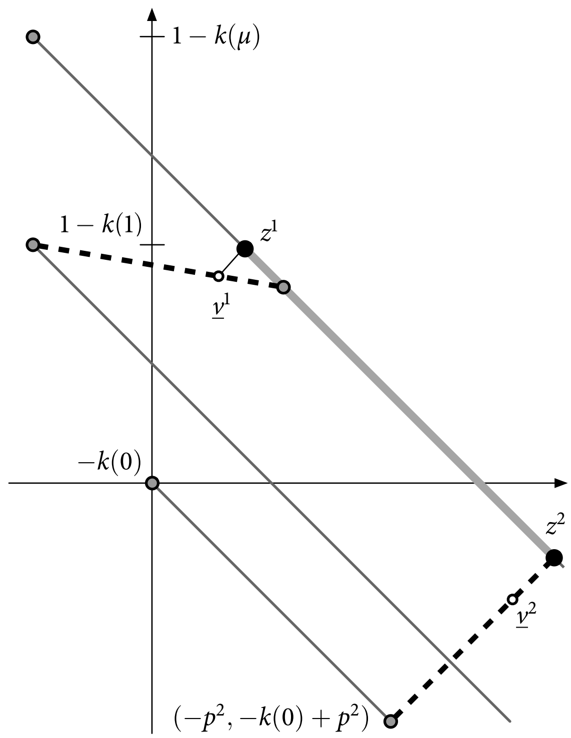

FIGURE 5. CONTRACTUAL EQUILIBRIUM WITH MANAGER OPTION.
Figures in this subsection are drawn to scale using the same parameters as in Figure 1.

the span $d = z_1^2 - z_1^1$ is largest for payments that satisfy:

$$p^1 - k(\mu^1) = p^2 - k(\mu^2), \tag{17}$$

and therefore the span is

$$d = (k(\mu^1) - k(\mu^2))(1 - \pi_1) - (\beta/\mu^1 - \beta) + \pi_1(1 - \beta). \tag{18}$$

We see that the maximal span, $(1 - \pi_1)(k(1) - k(0)) + \pi_1(1 - \beta)$, is attained when μ^1 is maximal and μ^2 minimal, and so the stationary part of the optimal external contract sets $\mu^1 = 1$, $\mu^2 = 0$, and payments to satisfy Equation 17. The resulting span is larger than the maximal span without option contracts, which was $\pi_1(1 - \beta)$. Continuation values in the contractual equilibrium are displayed in Figure 5.

The option contract provides flexibility to adjust the monitoring level under disagreement such that the cost is inefficiently high only in the case where shirking is to be avoided

obtain continuation value z^2 from the start of the next period.

($\mu^1 = 1$). This results in a larger span than what is obtained without the option contract, where the monitoring level is inefficiently costly in all cases under disagreement.

While the long-term option contract specifies either maximal or no monitoring, in equilibrium the parties agree each period to an intermediate monitoring level $\hat{\mu}$, which is the minimal level necessary to induce effort from the worker. As in the analysis of Section 1.2, this is the monitoring level for which the required high-signal bonus ρ just fits within the span d; and is thus given by $\rho = \frac{1-\delta}{\delta}\frac{\beta}{\mu} = d$. The larger span allows for a lower $\hat{\mu}$ on the equilibrium path and thereby a higher level of welfare.

Allocation of decision rights

The right to select an option can in principle be contracted on and externally enforced. Above we assumed that this right belongs to the manager. Suppose instead that the worker has the right.[19] Consider again options of the form (p^i, μ^i), $i = 1, 2$, where (p^i, μ^i) is intended to be selected under disagreement to punish player i, and p^i is a payment from the worker to the manager.

The disagreement value that is worst for the worker now entails her selecting option (p^1, μ^1) and then exerting effort; the disagreement value that is worst for the manager entails the worker selecting option (p^2, μ^2) and then shirking. The requirements of incentive compatibility and internal bargaining consistency then lead to the following span for the continuation values (see Appendix B.3):

$$d = \pi_1(1 - \beta - k(\mu^1) + k(\mu^2)) \tag{19}$$

The payoffs are again given by Equation 16, but the binding constraint for the worker to choose the desired option now requires that the worker's immediate payoff from selecting option (p^1, μ^1) and working must be no less than her immediate payoff from selecting the other option and shirking. This leads to the constraint $(\beta/\mu^1 - \beta) - p^1 \geq -p^2$.

The resulting span given in Equation 19 is increasing in μ^2 and decreasing in μ^1. The latter monitoring level, which under disagreement is used to punish the worker, cannot be too small because the high-signal bonus $\rho = \frac{1-\delta}{\delta}\frac{\beta}{\mu^1}$ required to induce high effort cannot exceed d. From this we see that the maximal span is obtained when $\mu^2 = 1$ and μ^1 is the smallest solution to

$$\frac{1-\delta}{\delta} \cdot \frac{\beta}{\mu^1} = d = \pi_1(1 - \beta - k(\mu^1) + k(1)). \tag{20}$$

[19] For example, a traveling salesman or service worker may have the right to control the extent to which his movements are registered by a global positioning system.

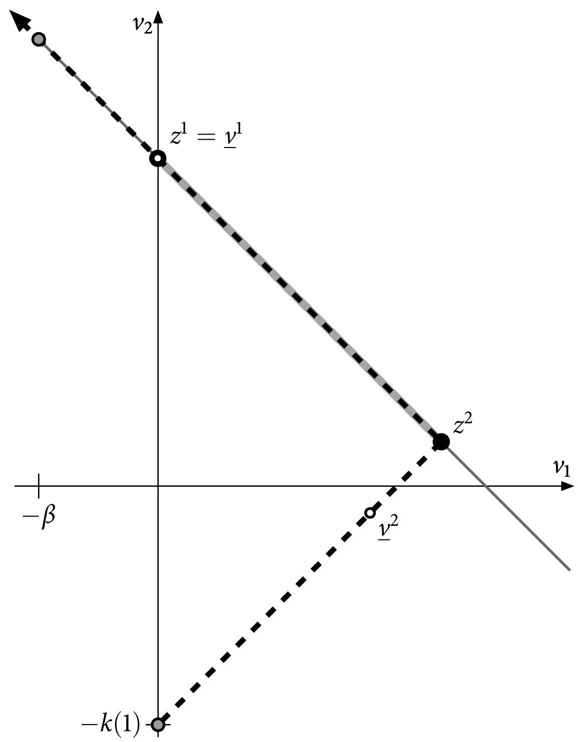

FIGURE 6. CONTRACTUAL EQUILIBRIUM WITH WORKER OPTION. The stage game payoff $(-\beta - p^1, 1 - k(\mu^1) + p^1)$ obtained when the worker is being punished is outside the range of the graph, in the direction of the heavy arrow.

Compared to the case of no option contract, where the span is $\pi_1(1-\beta)$, the span is here clearly larger. As in the case of manager control of the option, this occurs because the option allows the two monitoring levels μ^1 and μ^2 to be different.

Finally, to implement effort under agreement, the agreed upon monitoring level $\hat{\mu}$ must satisfy $\frac{1-\delta}{\delta}\frac{\beta}{\hat{\mu}} \leq d$, where d is given by Equation 20. This implies $\hat{\mu} = \mu^1$, i.e. that the monitoring levels under agreement and under disagreement to punish the worker should be equal. Consequently, in this case of worker control, the externally enforced terms in the inherited contract need not be renegotiated in the current period. Keeping the option menu fixed, the worker has incentives to select the appropriate monitoring intensity μ_i under disagreement, and the parties will agree on $\hat{\mu} = \mu^1$ otherwise. Contractual equilibrium values are shown in Figure 6.

The external contract in this case can be seen as specifying a "normal" level of monitoring to be applied whenever the worker is supposed to provide effort, and a very costly level

($\mu^2 = 1$) when it is intended that she shirks (under disagreement).[20] Note that the costly monitoring is not used to provide an effort incentive to the worker; rather it is used to "burn money" under disagreement. This arrangement shifts the disagreement point in a way that relatively favors the worker under disagreement when the worker is being rewarded.

Recall that in the case of manager control, the largest feasible span is obtained with inefficiently high monitoring when the worker is supposed to provide effort under disagreement ($\mu^1 = 1$). The difference between the two cases reflect the differences in the two parties' incentives when choosing between options. If the allocation of decision rights is negotiable—i.e., the external enforcer allows for both manager-option and worker-option contracts—then by semi-stationarity it is optimal for the parties to always specify the same stationary option contract in all future periods. That is, control rights under disagreement do not shift with the history. Comparing worker control with managerial control from the subsection above, we see that managerial control attains higher welfare when $\pi_2(k(1) - k(0)) \geq \pi_1(k(1) - k(\mu^1))$, while worker control attains higher welfare in the converse case. Managerial control is thus better if the manager's bargaining strength is sufficiently high relative to the monitoring cost function. But if the worker has sufficiently high bargaining power, then worker control is strictly better.

4.2 Multitasking

We next consider an example with a worker who exerts effort on two dimensions, both visible to the worker and the manager but only one verifiable to the external enforcer. For concreteness we refer to the verifiable dimension as "quantity" (a_2) and the non-verifiable dimension as "quality" (a_1). The good belongs to the manager. Payments conditional on quantities of the good can be externally enforced. The worker has costs $\kappa(a_1, a_2)$, and the manager's gross value is $\nu(a_1, a_2)$. Suppose quality is binary—either high ($a_1 = a_{1h}$) or low ($a_1 = 0$)—and that a high quality good is more costly for the worker and more valuable for the manager, so $\nu(a_{1h}, a_2) - \kappa(a_{1h}, a_2) > \nu(0, a_2) - \kappa(0, a_2) \geq 0$ for all quantities $a_2 \in (0, \bar{a}_2)$. Assume also that $\nu(a_1, 0) = \kappa(a_1, 0) = 0$ for both qualities.

The interesting problem entails how to attain high quality from the worker, since the optimal quantity of the low-quality good can be implemented by an external contract. Let z^1 and z^2 be the endpoints of the contractual-equilibrium value set, where as usual z^1 is worst for the worker and z^2 is worst for the manager. The worker can be given the incentive

[20] The level of inefficiency may be limited by constraints neglected here, such as participation constraints for the manager. This can be included in the analysis in a straightforward manner.

to choose high quality if

$$(1-\delta)\left(p-\kappa(a_{1h},a_2)\right)+\delta z_1^2 \geq (1-\delta)\left(p-\kappa(0,a_2)\right)+\delta z_1^1,$$

which simplifies to

$$\kappa(a_{1h},a_2)-\kappa(0,a_2) \leq \frac{\delta}{1-\delta}(z_1^2-z_1^1). \tag{21}$$

The welfare level attained is $L=\nu(a_{1h},a_2)-\kappa(a_{1h},a_2)$. The condition makes clear that the span must be sufficiently large to make up for the worker's additional cost to provide high rather than low quality; and that the higher the span, the higher is the welfare level that can be sustained. For the interesting case where the first best cannot be sustained, welfare will increase with a larger quantity of the high quality good.

We show that if the worker's bargaining power is sufficiently strong then the largest equilibrium span is obtained by letting the externally enforced terms take the form of a payment schedule that exactly compensates the worker for the costs to produce low quality, conditional on the verifiable quantity. Moreover, these terms need not be renegotiated in equilibrium in any period, so the external contract can be taken to be fully stationary in this application. In this payment scheme, the worker can be seen as having the right to choose a quantity-payment pair from the schedule under agreement, and, more importantly, under disagreement. In contrast, if the worker's bargaining power is weak, then it is instead optimal to give the manager the right to select quantities and payments from a menu. In this case the externally-enforced payment equals the manager's value for low quality at the verifiable provided quantity.

Details of the analysis are given in Appendix B.4. Since there are externally enforced transfers that can be made conditional on verifiable quantities, the feasible options are here a set of quantity-payment pairs. Two options are sufficient for disagreement play, specifying quantity-payment pairs (a_2^i, p^i), $i=1,2$, intended to be selected when punishing the worker and punishing the manager, respectively. The payments are from the manager to the worker. Incentive compatibility for the worker's selection is ensured by letting the options take the form of cost reimbursements for the low quality of the quantity supplied, i.e. $p^i=\kappa(0,a_2^i)$, $i=1,2$.

Under disagreement when being punished, the worker supplies high quality at the largest quantity that can be implemented.[21] Since he is reimbursed via the option contract for the cost of low quality $\kappa(0,a_2)$, this is the same quantity a_2 as given by Inequality 21.

[21] Any deviation by the worker is punished by coordination on his worst continuation equilibrium z^1, and selection of the desired option is then ensured if $p^1-\kappa(0,a_2^1)\geq p^2-\kappa(0,a_2^2)$. This constraint binds in equilibrium, and leads to payments in the form of cost reimbursements for low quality ($p^i=\kappa(0,a_2^i)$).

Hence welfare is L as under agreement and there is no renegotiation gain. The worker, just compensated for his costs (in part from the externally enforced payment and in part from internal rewards as given in Inequality 21), is left with a payoff of zero, and thus $z_1^1 = 0$.

Under disagreement when being rewarded, the worker supplies zero quantity and quality, yielding welfare level zero. Such an outcome is avoided in equilibrium by negotiations where the parties agree to act so as to achieve equilibrium welfare L, and to split the gains relative to disagreement play in accordance with their bargaining weights. The worst continuation value for the manager, and thus best for the worker, then yields $z_1^2 = \pi_1 L$ to the worker.

The resulting span is $d = z_1^2 - z_1^1 = \pi_1 L$, and from Inequality 21 it follows that the equilibrium quantity of the high quality good is the largest solution to

$$\kappa(a_{1h}, a_2) - \kappa(0, a_2) = \frac{\delta}{1-\delta}\pi_1(\nu(a_{1h}, a_2) - \kappa(a_{1h}, a_2)). \tag{22}$$

The external contract can be specified as a payment schedule $p(a_2') = \kappa(0, a_2')$, for any $a_2' \geq 0$, i.e., specifying that the worker is paid the cost of supplying the low quality good. It is then incentive compatible for the worker to provide $a = (a_{1h}, a_2)$, not only under agreement, but also under disagreement when being punisehd; and to provide $a = (0, 0)$ under disagreement when being rewarded. The external contract thus need not be renegotiated, and can be fully stationary.

It is of interest to compare this setting to an environment with no external enforcement. It follows from Miller and Watson (2013) that the contractual equilibrium in that case is to provide the high-quality good in a quantity given as the largest solution to an equation identical to Equation 22, except with the single term $\kappa(a_{1h}, a_2)$ on the left-hand side.[22] External enforcement thus improves welfare by allowing a larger quantity of the high quality good to be sustained in equilibrium.

The right to select the externally enforced terms may alternatively be allocated to the manager, thus giving her the right to select from the two externally enforced options (a_2^i, p^i), $i = 1, 2$. As the external enforcer is able to compel specific performance, the worker must comply with the specified quantity when the manager selects (a_2^i, p^i).[23] The analysis reveals that the maximal span in the case of manager control is obtained with options specifying payments $p^i = \nu(0, a_2^i)$, $i = 1, 2$. The external contract thus requires the manager to

[22] Internal incentives must then be sufficiently large to make up for the worker's total cost to provide high quality, not only for the cost differential as in Equation 22.

[23] If specific performance is not enforced, we may assume that deviations from the specified quantity are discouraged by sufficiently low associated externally enforced punishments.

pay her low-quality valuation of the supplied quantity.[24] These payments ensure incentive compatible selection on the part of the manager. The maximal span (and hence welfare) in this case depends on the two parties' bargaining powers, and exceeds the span under worker control when the manager's power (π_2) is sufficiently high. The comparison between the two schemes, manager control and worker control, reveals that the optimal allocation of decision rights in this model depends on the parties' relative bargaining power, and hence on the institutional context underlying their relationship.

4.3 Partnership

Consider a partnership setting that generalizes the Prisoners' Dilemma example from Miller and Watson (2013). The partners have equal bargaining power (i.e., $\pi_1 = \pi_2 = \frac{1}{2}$), and each partner i either exerts high effort (i.e., plays $a_i = 1$), or low effort (i.e., plays $a_i = 0$). Suppose the contracting setting has no verifiable information, but the external enforcer can impose a production technology $T \in [0, 1]$. Given a technology T, the stage game is the following convex combination of stage games:

	$a_2 = 1$	$a_2 = 0$
$a_1 = 1$	$1-\beta, 1-\beta$	$-\sigma, 1$
$a_1 = 0$	$1, -\sigma$	$0, 0$

with weight $1 - T$

	$a_2 = 1$	$a_2 = 0$
$a_1 = 1$	$\xi-\beta, \xi-\beta$	$0, \xi-\beta$
$a_1 = 0$	$\xi-\beta, 0$	$0, 0$

with weight T

where $0 < \beta < \xi < 1$ and $\sigma > 1 + \beta - \xi$.

With technology $T = 0$, this stage game is a "triangular" Prisoners' Dilemma (joint payoffs when one partner shirks while the other works are lower than when both shirk), such that each partner's payoff when both partners work is $1 - \beta$. If the external enforcer is willing to impose only $T = 0$, Miller and Watson (2013) show that under these assumptions the unique contractual equilibrium outcome is for both partners to exert low effort forever, regardless of how patient they are. The problem is that asymmetric play under disagreement must involve action profiles $a = (0, 1)$ and $a = (1, 0)$, which when $T = 0$ are so expensive to enforce that the necessary span of continuation values cannot be supported by the disagreement points they generate.

Alternative technologies $T > 0$ are "safer" in the sense that they improve the payoff for a partner who exerts high effort while his partner shirks, but at the cost of reducing the

[24] In particular, under disagreement to punish the manager, the regime calls for the worker to supply low quality, and for the parties to coordinate on the manager's worst continuation equilibrium z^2 for any outcome. Selection of the desired option is then incentive compatible if $\nu(0, a_2^2) - p^2 \geq \nu(0, a_2^1) - p^1$. This constraint binds in equilibrium, and leads to payments $p^i = \nu(0, a_2^i)$.

payoffs if both partners exert high effort. Indeed, for $T = 1$ this game lies at the boundary between a Prisoners' Dilemma and Chicken, such that each partner's payoff when both partners work is $\xi - \beta$.

Such a technology might operate by improving the productivity of individual effort, but requiring extensive and costly coordination to be effective when employed by both partners.[25] We will see that the partners can obtain efficient payoffs if they are patient enough, by specifying that they should use the safest technology under disagreement.

If the external enforcer is willing to impose any $T \in [0, 1]$ the partners write into their external contract, then we show that this capability enables the partners to obtain efficient payoffs if they are patient enough. First we construct an incentive compatible and internally bargain-consistent equilibrium in which the partners can support mutual cooperation $a = (1, 1)$ under $T = 0$ along the equilibrium path if they are patient enough. Let the external contract terms they write in each period specify $T^* = 0$ for the current period, with technology $\hat{T} = 1$ for all future periods. That is, under agreement they play a triangular Prisoners' Dilemma, but under disagreement they play Chicken. They cooperate along the equilibrium path, so the level is $L = 2(1 - \beta)$.

Consider a history off the equilibrium path, when partner 1 is supposed to be punished and the partners have just disagreed. Then they play a Chicken game with $\hat{T} = 1$ in the current period, and expect to renegotiate in the following period. Since $a = (1, 0)$ is a stage game equilibrium in Chicken, it can be supported by a continuation value that does not change with the stage game outcome. In this case the regime specifies $a = (1, 0)$ followed by player 1's worst continuation payoff vector, z^1. This plan generates a disagreement continuation payoff vector of

$$\underline{v}^1 = \big(\delta z_1^1, (1 - \delta)(\xi - \beta) + \delta z_2^1\big). \tag{23}$$

Now step back to the start of the period. Knowing that $\underline{v}^1$ is what they will get if they disagree, they renegotiate to the payoff vector z^1, characterized by:

$$z^1 = \underline{v}^1 + \frac{1}{2}\big(L - \underline{v}_1^1 - \underline{v}_2^1, L - \underline{v}_1^1 - \underline{v}_2^1\big) \tag{24}$$

For $a = (1, 1)$ to be incentive compatible when they play the $T^* = 0$ Prisoners' Dilemma

[25] Alternatively, the model may represent two parties that are affected by pollution, but where each may do costly abatement activity. The abatement benefits are equally shared between the parties, irrespective of who undertakes abatement. Technology T may then be thought of as reducing the cost of single-party abatement at the expense of increasing the cost of joint abatement, e.g. because the technology requires a scarce input that is more costly to use if more than one party uses it.

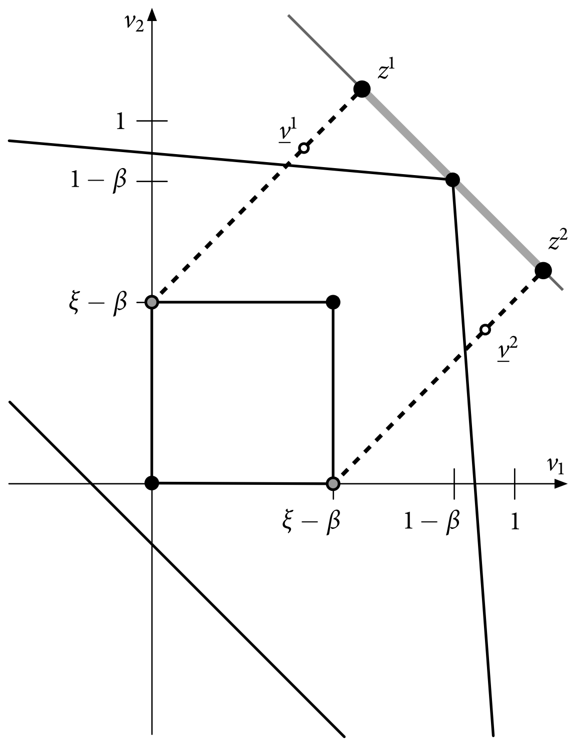

FIGURE 7. CONTRACTUAL EQUILIBRIUM IN PARTNERSHIP.

on the equilibrium path, it must be that $\frac{1}{2}L \geq (1-\delta) + \delta z_1^1$. This is the case if

$$\delta \geq \frac{2\beta}{\beta+\xi} < 1. \tag{25}$$

Intuitively, the Chicken game they play after a disagreement has two pure strategy stage game equilibria that are Pareto incomparable (among others). If a disagreement occurs, they coordinate to play the equilibrium that is relatively worse for whichever player is being punished. Knowing that this is what will happen if a disagreement occurs, when renegotiating they agree on a point in the equilibrium value set that is relatively worse for the player who is being punished; they implement this point by first having that player pay a transfer, and then continuing with efficient equilibrium path play of $a = (1,1)$ in the stage game with $T = 0$. The size of the transfer is constant in average utility terms, which means the size of the transfer in total utility terms can be made arbitrarily large by choosing a sufficiently high discount factor. When they are patient enough the transfer suffices to enforce efficient equilibrium path play. The contractual equilibrium is displayed in Figure 7

By Theorem 3, the contractual equilibrium is semistationary. We show in Supplementary Appendic B.5 that, for the stationary part of the external contract, in fact $\hat{T} = 1$ is superior to any $\hat{T} < 1$. Intuitively, asymmetric play in Chicken requires no incentives, allowing the continuation value to simply be z^1 under disagreement when partner 1 is being punished. In contrast, asymmetric play in any Prisoners' Dilemma with $\hat{T} < 1$ requires incentives, so the continuation value must give partner 1 more than z_1^1 for playing $a_1 = 1$ when being punished under disagreement.

5 Related Literature

The analysis of relational contracts was initiated by Klein and Leffler (1981), Shapiro and Stiglitz (1984), Bull (1987) and MacLeod and Malcomson (1989).[26] Levin (2003) showed that with transfers, optimal self-enforced contracts in a time invariant environment can be taken to be stationary. Levin also observed that optimal stationary contracts are "strongly optimal" in the sense that, for any feasible history, the continuation contract onwards is optimal. This variant of renegotiation proofness was further pursued by Goldlücke and Kranz (2013). With transfers, perfect monitoring, and no external enforcement, they show that Pareto-optimal subgame perfect payoffs and "strongly optimal" payoffs can generally be found by restricting attention to a simple class of stationary contracts.

Relative to renegotiation proofness, contractual equilibrium entails a different approach to equilibrium selection. The contrasts are discussed in depth in Miller and Watson (2013). Suffice it here to point out that, unlike contractual equilibrium, renegotiation proofness rules out renegotiation rather than modeling it explicitly, and thus does not account for the possibility of disagreement. Safronov and Strulovici (2016) also model renegotiation explicitly and allow for disagreements in a repeated game setting but without external enforcement. Their approach to bargaining is more permissive, allowing players to be punished for proposing Pareto improvements, and hence their solution concept makes substantially less sharp predictions than does contractual equilibrium.

Apart from the contractual equilibrium framework, this paper departs from previous analyses of relational contracting with some elements of external enforcement by allowing the external contract to make arbitrary long-term prescriptions, whereby the specifications in a given period are a function of the verifiable history of productive actions, restricted only by the capabilities of the external enforcer. This has important consequences, including that

[26] While the formal literature starts with Klein and Leffler, the concept of relational contracts had was first defined and explored by legal scholars (e.g., Macaulay 1963; Macneil 1978).

optimal external contracts are non-stationary in relevant settings.

The literature has shown that optimal relational contracts in time invariant environments may be non-stationary due to up-front bargaining and commitment to a long-term contract by one party (the principal) as in Ray (2002), limited liability (Fong and Li 2017), or persistent private information and limited enforcement (Martimort, Semenov, and Stole 2016). No such features are present in the model analyzed here; rather we show that limited external enforcement alone may make the equilibrium external contract non-stationary. We find that the optimal (semi-stationary) external contract specifies the same externally enforced terms for every future period but special terms for the current period. As we have noted, similar features arise in the complementary model of Kostadinov (2017).

On the general theme of choosing external contracts to operate in concert with self-enforcement, Iossa and Spagnolo (2011) have pointed out that it is common practice for contracting parties to write external contracts that contain inefficient clauses, but where these clauses are ignored in equilibrium. They explain this practice by observing that an inefficient external contract can be used as a credible threat to sustain a more efficient regime. Bernheim and Whinston (1998) emphasize that, when some aspects of performance are unverifiable, it is often optimal to leave other verifiable aspects of performance unspecified, so that optimal contracts are "less complete" than they could have been.[27] In a contractual equilibrium, the optimal external contract may entail such flexibility. We have noted that it takes the form of options in some applications, and that allocation of decision rights can be important is such settings.

Baker, Gibbons, and Murphy (2011) also demonstrate how allocation of such rights matters in relational contracting, but via a channel very different from ours. They analyze how governance structures (allocations of control) can facilitate relational contracts that improve on spot transactions in settings where such transactions would produce inefficient adaptation to changing circumstances. Relatedly, Barron, Gibbons, Gil, and Murphy (2015) analyze self-enforced agreements that facilitate efficient adaptation (they call it relational adaptation), and show how these agreements, combined with an external contract, induce state-dependent decision-making that improves upon the expected payoffs under either external contracting or relational contracting alone. Their theoretical model assumes stationarity of equilibrium strategies and Nash reversion (permanent punishment following

[27] Iossa and Spagnolo (2011) examine a repeated principal-agent model in which, in each period, players have the option to trigger penalties specified by the external contract. Long-term external contracts are restricted to be stationary. Renegotiation is costly and disagreement results in adherence to an inefficient external contract in all future periods. Bernheim and Whinston (1998) examine a class of two-period contracting problems with both external enforcement and self-enforcement.

any deviation).

Finally, a considerable literature has investigated the implications of renegotiation and the "hold-up problem" in short-term trading relationships in which unverifiable investments are followed by renegotiation and then verifiable trade.[28] Researchers have shown that the hold-up problem can be alleviated in some short-term trading relationships, in particular in settings of "own-investment" (such as in Aghion, Dewatripont, and Rey 1994, Noldeke and Schmidt 1995, and Edlin and Reichelstein 1996). Results in this literature rely on complementarities between productive actions. Specifically, investment decisions influence the value of trade. In our model, as with most in the relational-contracting literature, all productive actions occur at one instant in each period, meaning that production and delivery are integrated or simultaneous. Thus, the conditions for achieving efficiency that are developed in the hold-up literature are not present here. It would be interesting in future work to examine settings with technological state variables, where the productive actions taken in one period influence the payoffs received in future periods.

6 Conclusion

In the words of Malcomson (2013), "The literature on relational contracts is concerned with the impact of the on-going nature of the relationship on trade between the parties, on their payoffs, on the nature of any legally enforceable contract that is used to supplement the relational contract, and on the design of organizations." This paper has focused on the nature of the externally enforced part of the contract, and on its implications for the overall relationship. We have modeled this in a general framework allowing for various forms of external enforcement, by extending the concept of contractual equilibrium for infinite horizon games in Miller and Watson (2013) to such environments. In a contractual equilibrium the parties can re-evaluate and renegotiate all aspects of their relationship each period, including their externally enforced contract and their regime.

In contrast to most previous analyses, we have allowed the externally enforced part of the contract to be general and restricted only by verifiability. This opens the possibility of endogenous non-stationarity, and we have showed that in interesting and relevant environments the equilibrium long-term external contract is indeed non-stationary. When

[28] Prominent entries include Hart and Moore (1988), Hart and Moore (1999), Noldeke and Schmidt (1995), Che and Hausch (1999), Segal (1999), and Maskin and Tirole (1999); see Bolton and Dewatripont (2005) for a survey. Most closely related are models with individual trade actions, such as Watson (2007), Evans (2008), and Buzard and Watson (2012). Because our theory treats renegotiation explicitly and incorporates bargaining power, negotiations in a contractual equilibrium operate similarly to what is explored in the hold-up literature.

the external enforcer can compel monetary transfers, this contract is semi-stationary: the long-term part which governs future periods is stationary, but the provisions for the current period are different. In every period, the terms for this period are renegotiated to achieve the highest joint value that can be attained.

In addition to laying out the general framework and proving existence and properties of contractual equilibrium, we have explored some applications of the theory. A monitoring case illustrates the nature of semi-stationarity in a very simple setting. An extended version also illustrates the value of strategic flexibility, implemented as options, and the importance of decision rights. An application to a principal-agent environment with multitasking further shows how the external contract can take the form of simple payment schedules for the agent. An application to partnership with moral hazard demonstrates that the long-term part of the external contract may feature highly inefficient terms that make the partners' efforts less complementary.

A Recursive Characterization and Proofs

In this appendix, we perform analysis that yields a recursive characterization of contractual equilibrium payoffs, along the lines of Abreu, Pearce, and Stacchetti (1990) and Miller and Watson (2013), where one relates continuation values that can be achieved from a given period to the continuation values in the next period. The key complication we face here is that the set of continuation values generally differs across periods and must be indexed by the inherited external contract. Thus, instead of looking for a fixed point set of continuation values, as is the case in the earlier literature, we are looking for a fixed point in the space of *indexed collections of* sets of continuation values.

The first subsection formulates self-generation of the continuation-value sets. The second subsection shows how the analysis can be done with a normalized version of the continuation values and gives an existence condition. The third subsection contains a proof of Theorem 2. [29]

A.1 Self-Generation

To get into the details, let us start by describing continuation values that can be achieved from the action phase in a given period t. For the following definition, take as given any collection $\mathcal{W} = \{W(c')\}_{c' \in C}$, where $W(c') \subset \mathbb{R}^2$ for every $c' \in C$. This collection

[29] An additional existence result (for settings with a finite number of external contracts and finite stage games) is provided in Appendix B.1.

describes the continuation values that can be selected from the start of period $t+1$, as a function of the external contract inherited in period $t+1$. We allow $W(c') = \emptyset$.

As defined in Section 3.2, for any $\gamma = (A, X, \lambda, u, P) \in G$ and $z\colon X \times [0,1] \to \mathbb{R}^2$, let $\overline{u}^\gamma(a) \equiv E_x[\, u(a,x)\,]$ and let $\overline{z}^\gamma(a) \equiv E_{x,\phi}[\, z(x,\phi)\,]$, where the first expectation is taken with respect to $x \sim \lambda(a)$ and the second expectation is taken with respect to $x \sim \lambda(a)$ and $\phi \sim U[0,1]$.

Definition 10. *Take any $c \in C$ and let $g(c) = (A, X, \lambda, u, P)$ be the stage game designated for the current period under external contract c. Say that $w \in \mathbb{R}^2$ is* ***c-supported*** *relative to $\mathcal{W}$ if there exists a mixed action profile $\alpha \in \Delta A$ and a function $z\colon X \times [0,1] \to \mathbb{R}^2$, such that*

- *for all $x \in X$ and $\phi \in [0,1]$, it is the case that $z(x,\phi) \in W(\zeta(c,x,\phi))$;*
- *α is a Nash equilibrium of $\langle A, (1-\delta)\overline{u}^{g(c)} + \delta \overline{z}^\gamma \rangle$; and*
- *the expected payoff of this Nash equilibrium is w.*

In reference to the first two conditions, we say that α is ***c-enforced*** *relative to $\mathcal{W}$.*

Suppose that the players enter a period with c as their inherited external contract, a particular continuation value $\underline{w}$ is c-supported, and the players will coordinate to achieve $\underline{w}$ in the event that they fail to reach an agreement in the negotiation phase. Further, suppose that the players are able to achieve a joint value of L through negotiation. Then, from internal bargain consistency, the continuation value will be $w = \underline{w} + \pi(L - \underline{w}_1 - \underline{w}_2)$ from the beginning of the period, presuming $L \geq \underline{w}_1 + \underline{w}_2$.

For a given collection $\mathcal{W} = \{W(c')\}_{c' \in C}$, any $c \in C$, and any level L, let $B^L(c, \mathcal{W})$ be the set of continuation values that can be achieved from the beginning of a period in which c is the inherited contract and the players can bargain to obtain level L:

$$B^L(c, \mathcal{W}) = \{\underline{w} + \pi(L - \underline{w}_1 - \underline{w}_2) \mid \underline{w} \text{ is } c\text{-supported relative to } \mathcal{W} \text{ and } L \geq \underline{w}_1 + \underline{w}_2\}. \tag{26}$$

We build in the condition $L \geq \underline{w}_1 + \underline{w}_2$ to identify just those values of $\underline{w}$ that are consistent with the level being L. Otherwise, some disagreement points could yield a strictly higher joint value and B^L would incorrectly represent that the players negotiate to achieve a lower level. Note that $B^L(c, \mathcal{W})$ may be empty.

Observe that, in a scenario in which L is the level, if $\mathcal{W}$ is the collection of continuation values available from the start of some period $t+1$, then $\{B^L(c', \mathcal{W})\}_{c' \in C}$ is the set of continuation values attainable from the start of period t.

Definition 11. *Consider a collection $\mathcal{W} = \{W(c')\}_{c' \in C}$. We say that $\mathcal{W}$ is **self-generating with level** $\mathbf{L}$ if there is an external contract $\overline{c} \in C$ and a $\overline{c}$-supported continuation value $\overline{w}$ that satisfies $L = \overline{w}_1 + \overline{w}_2$, and*

$$W(c) \subset B^L(c, \mathcal{W}) \text{ for all } c \in C. \tag{27}$$

Let us next apply operator B^L to characterize the continuation values for a single regime. Given a regime r, let $C(r) \equiv \{c \in C \mid c = \hat{c}(h) \text{ for some } h \in H(r)\}$. This is the set of external contracts that arise as inherited contracts for histories in $H(r)$. Then for every $c \in C(r)$, let $V(c; r)$ be the set of continuation values for regime r following histories in which c is the inherited external contract. That is,

$$V(c; r) = \{v(h; r) \mid h \in H(r), \hat{c}(h) = c\}.$$

Also, let us specify $V(c; r) = \emptyset$ for every $c \in C \setminus C(r)$. Finally, let $\mathcal{V}(r) \equiv \{V(c'; r)\}_{c' \in C}$.

Lemma 3. *If regime r is incentive compatible in the action phase and internally bargain-consistent, then $\mathcal{V}(r)$ is self-generating with level L, where L is the level of regime r.*

Proof of Lemma 3. Consider any $c \in C(r)$ and any $w \in V(c; r)$. By definition of $V(c; r)$ there is a history $h \in H(r)$ for which $v(h; r) = w$. By the condition of internal bargain-consistency and Lemma 1, we know that $v_1(h; r) + v_2(h; r) \geq \underline{v}_1(h; r) + \underline{v}_2(h; r)$ and

$$\begin{aligned} v(h; r) &= (1 - \delta) r^{\mathrm{m}}(h) + \tilde{v}(h, r^{\mathrm{c}}(h), r^{\mathrm{m}}(h), r^{\mathrm{a}}(h, r^{\mathrm{c}}(h), r^{\mathrm{m}}(h)); r) && (28) \\ &= \underline{v}(h; r) + \pi \left(v_1(h; r) + v_2(h; r) - \underline{v}_1(h; r) - \underline{v}_2(h; r) \right) && (29) \\ &= \underline{v}(h; r) + \pi (L - \underline{v}_1(h; r) - \underline{v}_2(h; r)). && (30) \end{aligned}$$

Let $\underline{w} = \tilde{v}(h, \hat{c}(h), 0, r^{\mathrm{a}}(h, \hat{c}(h), 0); r)$ and $\overline{w} = \tilde{v}(h, r^{\mathrm{c}}(h), r^{\mathrm{m}}(h), r^{\mathrm{a}}(h, r^{\mathrm{c}}(h), r^{\mathrm{m}}(h)); r)$. Individual rationality in the action phase implies that $\overline{w}$ is $r^{\mathrm{c}}(h)$-supported relative to $\mathcal{V}(r)$ and $\underline{w}$ is c-supported relative to $\mathcal{V}(r)$. By construction, $L = \overline{w}_1 + \overline{w}_2$ and

$$v(h; r) = \underline{w} + \pi(L - \underline{w}_1 - \underline{w}_2),$$

These facts imply that $v(h; r) \in B^L(c, \mathcal{V}(r))$. Thus $V(c; r) \subset B^L(c, \mathcal{V}(r))$ for all $c \in C(r)$. The same condition holds trivially for all $c \in C \setminus C(r)$ due to $V(c; r) = \emptyset$ in this case. $\square$

As one would expect, the reverse implication also holds.

Lemma 4. *If a collection $\mathcal{W} = \{W(c')\}_{c' \in C}$ is self-generating with level L and satisfies $W(c^0) \neq \emptyset$, then there is a regime r that is incentive compatible in the action phase and internally bargain-consistent, has level L, and has the property that $V(c; r) \subset W(c)$ for all $c \in C$.*

Proof of Lemma 4. The result follows from standard arguments, along the lines of the construction detailed in Miller and Watson (2013). We construct the regime r by, for select histories, specifying the behavior identified in the self-generation conditions. Our construction will cover all of the histories in $H(r)$.

Start with the null history, h^0, note that $\hat{c}(h^0) = c^0$, and pick any element $w \in W(c^0)$ to be the equilibrium continuation value from the beginning of the game. From the self-generation conditions, $w \in B^L(c^0, \mathcal{W})$ and so we can find an external contract $\overline{c}$, a $\overline{c}$-supported (relative to $\mathcal{W}$) value $\overline{w}$, and a c^0-supported (relative to $\mathcal{W}$) disagreement value $\underline{w}$ such that $L = \overline{w}_1 + \overline{w}_2$ and $w = \underline{w} + \pi(L - \underline{w}_1 - \underline{w}_2)$.

Prescribe $r^{\mathrm{c}}(h^0) = \overline{c}$ and let $r^{\mathrm{m}}(h^0)$ to be the corresponding transfer that achieves w as the continuation value from the beginning of period 1 when $\overline{w}$ is the continuation value from the action phase, so that $w = (1 - \delta) r^{\mathrm{m}}(h^0) + \overline{w}$. Then prescribe $r^{\mathrm{a}}(h^0, c^0, 0)$ to be the mixed action α that is identified by self-generation to c^0-support $\underline{w}$. Likewise, prescribe $r^{\mathrm{a}}(h^0, r^{\mathrm{c}}(h^0), r^{\mathrm{m}}(h^0))$ to be the mixed action identified to $\overline{c}$-support $\overline{w}$. For other values of (c^1, m^1), the prescribed action profile $r^{\mathrm{a}}(h^0, c^1, m^1)$ can be arbitrary. Such a joint decision entails a joint deviation and thus the equilibrium conditions are not imposed for histories that follow it.

The construction continues by looking at all of the one-period histories in $H(r)$, which follow the joint actions specified in the previous paragraph (and all of the possible action profiles in $A(c^0)$ and $A(\overline{c})$). For each such history h, a specific continuation value from $W(\hat{c}(h))$ is required to provide the incentives and continuation payoffs specified in period 1. We simply repeat the steps in the previous paragraph to specify behavior in period 2 following history h. The process continues for period 3, 4, and so on, which inductively yields a fully specified regime. By construction from the self-generation conditions, the regime's continuation values have the desired properties and the regime is incentive compatible in the action phase and internally bargain-consistent. □

A.2 Normalization, Computation, and Existence

To review the analysis thus far, self-generation relates to the conditions of incentive compatibility in the action phase and internal bargain-consistency for a regime. To characterize contractual equilibrium, we must look across all regimes with these properties and find one with the highest level if it exists. Fortunately, we can compare regimes with different levels by normalizing to level zero and (for now) ignoring the condition that $L \geq \underline{w}_1 + \underline{w}_2$. To this end, for a vector $b \in \mathbb{R}^2$ and a collection of sets $\mathcal{W} = \{W(c)\}_{c \in C}$, let "$\mathcal{W} + b$" denote the collection that results by adding b to all the points in the sets. That is, $\mathcal{W} + b = \{W(c) + b\}_{c \in C}$, where $W(c) + b = \{w + b \mid w \in W(c)\}$.

Let us define a version of operator B^L that does not include the condition that $L \geq \underline{w}_1 + \underline{w}_2$. We then use the term *weak self-generation* as a modification of self-generation in which this inequality is left out and feasibility of L is not checked.

$$\hat{B}^L(c, \mathcal{W}) = \{\underline{w} + \pi(L - \underline{w}_1 - \underline{w}_2) \mid \underline{w} \text{ is } c\text{-supported relative to } \mathcal{W}\}. \tag{31}$$

Definition 12. *Consider a collection $\mathcal{W} = \{W(c')\}_{c' \in C}$. We say that $\mathcal{W}$ is* ***weakly self-generating with level L*** *if $W(c) \subset \hat{B}^L(c, \mathcal{W})$ for all $c \in C$.*

Note that the equilibrium incentive conditions in Definition 10, for the induced game $\langle\, A, (1-\delta)\overline{u}^{g(c)} + \delta \overline{z}^{\gamma} \,\rangle$, would not be affected by transforming all of the sets in $\mathcal{W}$ by a constant vector $b \in \mathbb{R}^2$. Transforming z by the same constant maintains feasibility and yields a strategically equivalent induced game $\langle\, A, (1-\delta)\overline{u}^{g(c)} + \delta \overline{z}^{\gamma} + \delta b \,\rangle$. The c-supported continuation values are all transformed by δb. Thus, w is ***c*-supported** relative to $\mathcal{W}$ if and only if $w + \delta b$ is ***c*-supported** relative to $\mathcal{W} + b$. Using the definition of $\hat{B}^L(c, \mathcal{W})$ in Equation 31, this means that $\hat{B}^L(c, \mathcal{W} + b) = \hat{B}^L(c, \mathcal{W}) + \delta b - \delta(b_1 + b_2)\pi$. We also see that changing the value of L causes the set $\hat{B}^L(c, \mathcal{W})$ to merely shift by a constant multiple of the vector π, so that $\hat{B}^L(c, \mathcal{W}) = \hat{B}^0(c, \mathcal{W}) + \pi L$.

These facts imply (letting $b = -\pi L$) that a collection $\mathcal{W}$ is weakly self generating with level L if and only if the collection $\mathcal{W}^0 \equiv \mathcal{W} - L\pi$ is weakly self generating with level 0. So, we can replace the condition on $\mathcal{W}$ in Definition 12 with the normalized version where $L = 0$ and the collection is $\mathcal{W}^0$. Let us say that any given collection $\mathcal{W}^0$ is **normalized self generating** if it satisfies $W^0(c) \subset \hat{B}^0(c, \mathcal{W}^0)$ for all $c \in C$.

Let $\mathcal{Y} = \{Y(c)\}_{c \in C}$ be the union of all normalized self-generating collections of continuation-value sets. That is, for every $c \in C$, $Y(c)$ is the union of the sets $W^0(c)$ across all normalized self-generating collections $\mathcal{W}^0 = \{W^0(c)\}_{c \in C}$. It is clear that operator $\hat{B}^0$ is monotone in $\mathcal{W}$ and the set of c-enforced action profiles is increasing in $\mathcal{W}$. Thus, $\mathcal{Y}$ is normalized self-generating and contains every other normalized self-generating collection. We have a sufficient condition for existence based on this union collection.

Lemma 5. *For any relational-contract setting, a sufficient condition for the existence of a contractual equilibrium is that $Y(c^0) \neq \emptyset$ and the following optimization problem has a* *solution:*

$$\max_{\substack{c \in C \\ \alpha \in \Delta A(c)}} \quad \overline{u}_1^{g(c)}(\alpha) + \overline{u}_2^{g(c)}(\alpha) \tag{32}$$
$$\textit{subject to:} \quad \alpha \text{ is } c\text{-enforced relative to } \mathcal{Y}.$$

The contractual-equilibrium level L^* is the maximized value of this optimization problem.

Proof of Lemma 5. Let L^* be the maximized value of Optimization Problem 32 and let $\overline{c}$ be a maximizer for this problem. Because $\mathcal{Y}$ is normalized self-generating, we know that $\mathcal{Y} + L^*\pi$ is weakly self-generating with level L^*.

We first show that $\mathcal{Y} + L^*\pi$ is self-generating (without the weak qualifier) with level L^*. Take any $c \in C$ and let w be any continuation value that can be c-supported relative to $\mathcal{Y} + L^*\pi$. We have that $w = (1-\delta)\tilde{u} + \delta\tilde{z}$, where $\tilde{u}$ is the expected payoff in a stage game and $\tilde{z}$ is the expected continuation value in the next period. We know $\tilde{z}_1 + \tilde{z}_2 = L^*$ and, by definition of L^*, we know that $\tilde{u}_1 + \tilde{u}_2 \leq L^*$, so we conclude that $w_1 + w_2 \leq L^*$. Recalling the definition of B^L in Expression 26, we see that the inequality $L^* \geq \underline{w}_1 + \underline{w}_2$ never binds in the construction of $B^{L^*}(c, \mathcal{Y} + L^*\pi)$, and so $B^{L^*}(c, \mathcal{Y} + L^*\pi) = \hat{B}^{L^*}(c, \mathcal{Y} + L^*\pi)$ for every $c \in C$. Thus, that $\mathcal{Y} + L^*\pi$ is weakly self-generating with level L^* implies that $\mathcal{Y} + L^*\pi$ is self-generating with level L^*.

Lemma 4 then establishes that there is a regime r^* that is incentive compatible in the action phase, is internally bargain-consistent, and has level L^*. To finish the proof, we need only show that no other incentive-compatible and internally bargain-consistent regime can have a level strictly greater than L^*. Consider any regime r with these properties and let $\mathcal{W}^0 = \{W^0(c')\}_{c' \in C}$ be its normalized collection of continuation-value sets. There is an external contract c such that the regime's level is achieved by a c-supported continuation value relative to $\mathcal{W}^0$. Since $W(c') \subset Y(c)$ for every $c' \in C$, the same level is attainable in Optimization Problem 32, and thus the level of regime r can be no greater than L^*. We have shown that r^* maximizes the level over all regimes that are incentive compatible in the action phase and internally bargain-consistent. ☐

A.3 Externally Enforced Transfers and Semi-stationarity

In this subsection, we provide a proof of Theorem 2, which is restated here:

Theorem 2. *Suppose Assumptions 1–3 hold and the contractual setting has externally enforced transfers. If $\Lambda(d)$ and $\Xi(d)$ exist for all $d \geq 0$ then there exists a semi-stationary contractual equilibrium. The contractual-equilibrium span d^* is the largest fixed point of Λ, which exists, and the level is $L^* = \Xi(d^*)$.*

One aspect of the proof involves establishing, for select external contracts $c \in C$, some facts about the span of the set $Y(c)$. Recall that the span of any set $W \subset \mathbb{R}^2$ of constant joint value is the vertical/horizontal distance between its extreme points:

$$\mathrm{Span}(W) \equiv \sup\{w_1 - w_1' \mid w, w' \in W\}.$$

We separate the proof into three steps.

Step 1: Λ has a maximal fixed point.

For a given distance d, let y^1, y^2, α^1, and α^2 solve Optimization Problem 14 to determine $\Lambda(d)$. Recall that $\omega(\alpha, \gamma, y)$ denotes the normalized continuation value if in the current period α is played in stage game γ and the continuation value in the next period is given by $y\colon X^\gamma \to \mathbb{R}^2_0$. This was defined in Subsection 3.2. From the definition of ω we have that

$$\begin{array}{rl} \Lambda(d) = & (1-\delta)(\pi_2 \overline{u}_1^\gamma(\alpha^2) - \pi_1 \overline{u}_2^\gamma(\alpha^2)) + \delta \overline{y}_1(\alpha^2) \\ & - \left[(1-\delta)(\pi_2 \overline{u}_1^\gamma(\alpha^1) - \pi_1 \overline{u}_2^\gamma(\alpha^1)) + \delta \overline{y}_1(\alpha^1)\right] \\ = & (1-\delta)\overline{u}_1^\gamma(\alpha^2) + \delta \overline{y}_1^2(\alpha^2) - \pi_1(1-\delta)(\overline{u}_1^\gamma(\alpha^2) + \overline{u}_2^\gamma(\alpha^2)) \\ & - \left[(1-\delta)\overline{u}_1^\gamma(\alpha^1) + \delta \overline{y}_1^1(\alpha^1)\right] + \pi_1(1-\delta)(\overline{u}_1^\gamma(\alpha^1) + \overline{u}_2^\gamma(\alpha^1)) \end{array}. \tag{33}$$

The following four inequalities, in order, follow from enforcement of α^1 (in particular that player 1 cannot gain by deviating to α_1^2), that the joint stage-game value exceeds $-\vartheta$ (from Assumption 3), enforcement of α^2 (in particular that player 2 cannot gain by deviating to α_2^1), and that the joint stage-game value is no greater than ϑ (from Assumption 3):

$$\begin{array}{rcl} -\left[(1-\delta)\overline{u}_1^\gamma(\alpha^1) + \delta \overline{y}_1^1(\alpha^1)\right] & \leq & -(1-\delta)\overline{u}_1^\gamma(\alpha_1^2, \alpha_2^1) - \delta \overline{y}_1^1(\alpha_1^2, \alpha_2^1) \\ 0 & \leq & (1-\delta)\overline{u}_1^\gamma(\alpha_1^2, \alpha_2^1) + (1-\delta)\overline{u}_2^\gamma(\alpha_1^2, \alpha_2^1) + (1-\delta)\vartheta \\ 0 & \leq & -(1-\delta)\overline{u}_2^\gamma(\alpha_1^2, \alpha_2^1) - \delta \overline{y}_2^2(\alpha_1^2, \alpha_2^1) \\ & & +(1-\delta)\overline{u}_2^\gamma(\alpha^2) + \delta \overline{y}_2^2(\alpha^2) \\ 0 & \leq & -(1-\delta)\overline{u}_2^\gamma(\alpha^2) - (1-\delta)\overline{u}_1^\gamma(\alpha^2) + (1-\delta)\vartheta. \end{array}$$

Summing these inequalities yields

$$\begin{array}{l} -\left[(1-\delta)\overline{u}_1^\gamma(\alpha^1) + \delta \overline{y}_1^1(\alpha^1)\right] \leq \\ \qquad - \delta \overline{y}_1^1(\alpha_1^2, \alpha_2^1) - \delta \overline{y}_2^2(\alpha_1^2, \alpha_2^1) + \delta \overline{y}_2^2(\alpha^2) - (1-\delta)\overline{u}_1^\gamma(\alpha^2) + 2(1-\delta)\vartheta. \end{array}$$

Substituting the bracketed left-side terms into Equation 33, simplifying, and using the fact that $\overline{y}_1^2(\alpha^2) + \overline{y}_2^2(\alpha^2) = 0$, we obtain

$$\Lambda(d) \leq 2(1+\pi_1)(1-\delta)\vartheta - \delta\overline{y}_1^1(\alpha_1^2, \alpha_2^1) - \delta\overline{y}_2^2(\alpha_1^2, \alpha_2^1).$$

Because $\overline{y}_1^1(\alpha_1^2, \alpha_2^1) \in [0, d]$ and $\overline{y}_2^2(\alpha_1^2, \alpha_2^1) \in [-d, 0]$, we conclude that

$$\Lambda(d) \leq 2(1+\pi_1)(1-\delta)\vartheta - \delta d. \tag{34}$$

In words, $\Lambda(d)$ is bounded above by a line with slope $\delta < 1$. We thus know that $\Lambda(d) < d$ for all $d > \overline{d}$ where $\overline{d}$ solves $\overline{d} = 2(1+\pi_1)(1-\delta)\vartheta - \delta\overline{d}$. Additionally, Λ is increasing and satisfies $\Lambda(0) \geq 0$. Thus, by Tarski's fixed-point theorem, Λ has a maximal fixed point, d^*. □

To simplify notation a bit, let us set the following notation that we will utilize in the next two steps of the proof. Let γ^*, y^{*1}, y^{*2}, α^{*1}, and α^{*2} denote any solution to Optimization Problem 14 for Λ evaluated at d^*. Let γ^{**}, y^{**}, and α^{**} denote any solution to Optimization Problem 15 for Ξ evaluated at d^*. Let $\gamma^0 = g(c^0)$ denote the stage game played if the players have failed to make an agreement to date. Further, let $u^0 \equiv u^{\gamma^0}$, $u^* \equiv u^{\gamma^*}$, and $u^{**} \equiv u^{\gamma^{**}}$.

Step 2: A stationary external contract achieves the largest span.

Let Z denote the set of continuation values achieved when we restrict the range of y to be $\mathbb{R}_0^2(d)$ for a given span d:

$$Z(\gamma, d) \equiv \{\omega(\alpha, \gamma, y) \mid y\colon X^\gamma \to \mathbb{R}_0^2(d) \text{ and } \alpha \text{ is enforced relative to } \gamma \text{ and } y\}.$$

Let $\hat{d} \equiv \sup\{\mathrm{Span}(Y(c)) \mid c \in C\}$.

Let us next review the technical steps presented in Subsection 3.2. There we presumed that the maximum span exists and is achieved, meaning there is an external contract $\tilde{c}$ and values $z^1, z^2 \in Y(\tilde{c})$ such that $z_1^2 - z_1^1 = \hat{d}$. Note that here we are expressing continuation values in their normalized version. In Subsection 3.2 we established that if w is c-supported relative to $\mathcal{Y}$ then there is an external contract c' that transitions always to $\tilde{c}$ such that w is c'-supported relative to $\mathcal{Y}$. That is, continuation values in the next period are all in $Y(\tilde{c})$. This implies that, for each $c \in C$, there is an external contract c' and a vector $b \in \mathbb{R}_0^2$ such that

$$Y(c) \subset Z(g(c'), \hat{d}) + b. \tag{35}$$

In fact, Relation 35 holds without having to assume that the maximum span exists and is achieved. To see this, note that the assumption was used only to identify the particular set $Y(\tilde{c})$ in which to place all continuation values. In converting to the Z formulation, we have substituted $\mathbb{R}^2_0(\hat{d})$ for $Y(\tilde{c})$, and by definition of $\hat{d}$ we know that for every $c'' \in C$ there is a vector b'' such that $Y(c'') \subset \mathbb{R}^2_0(\hat{d}) + b''$.

For every $c \in C$, Relation 35 implies that there exists a stage game $\gamma \in G$ such that $\mathrm{Span}(Y(c)) \leq \mathrm{Span}(Z(\gamma, \hat{d}))$. We also know that $\mathrm{Span}(Z(\gamma, \hat{d})) \leq \Lambda(\hat{d})$ because Optimization Problem 14 has the stage game as one of the choice variables. So we have that $\mathrm{Span}(Y(c)) \leq \Lambda(\hat{d})$ for every $c \in C$, which implies that $\sup_{c \in C} \mathrm{Span}(Y(c)) \leq \Lambda(\hat{d})$ and thus $\hat{d} \leq \Lambda(\hat{d})$. Because Λ is increasing and satisfies Inequality 34, we know the maximal fixed point satisfies $d^* \geq \hat{d}$.

Finally, we find a stationary external contract $c^* \in C$ for which $\mathrm{Span}(Y(c^*)) = d^*$. Let c^* be the external contract that specifies stage game γ^* and always transitions backs to itself. That is, $g(c^*) = \gamma^*$ and $\zeta(c^*, x, \phi) = c^*$ for all $x \in X^{\gamma^*}$ and $\phi \in [0, 1]$. Assumption 1 ensures that this external contract is an element of C. Now compare Optimization Problem 14 for d^* to the construction of $\hat{B}^0(c^*, \mathcal{W})$ in Equation 31, under the assumption that $W(c^*) = \{z^1, z^2\}$ where $z^1, z^2 \in \mathbb{R}^2_0$ satisfy $z^2_1 - z^1_1 = d^*$. In the latter, the players can select between z^1 and z^2 arbitrarily using the public draw ϕ, and thus achieve any continuation value in the convex hull, which is equivalent to the set $z^1 + \mathbb{R}^2_0(d^*)$. Thus, $\hat{B}^0(c^*, \mathcal{W})$ contains the two extreme points related to the solution of Optimization Problem 14:

$$
\begin{aligned}
z^{1\prime} &= (1-\delta)\left(\pi_2 \overline{u}^*_1(\alpha^{*1}) - \pi_1 \overline{u}^*_2(\alpha^{*1}),\ \pi_1 \overline{u}^*_2(\alpha^{*1}) - \pi_2 \overline{u}^*_1(\alpha^{*1})\right) + \delta(z^1 + \overline{y}^{*1}(\alpha^{*1})) \\
z^{2\prime} &= (1-\delta)\left(\pi_2 \overline{u}^*_1(\alpha^{*2}) - \pi_1 \overline{u}^*_2(\alpha^{*2}),\ \pi_1 \overline{u}^*_2(\alpha^{*2}) - \pi_2 \overline{u}^*_1(\alpha^{*2})\right) + \delta(z^1 + \overline{y}^{*2}(\alpha^{*2})).
\end{aligned}
$$

By construction, $z^{2\prime}_1 - z^{1\prime}_1 = d^*$. Recall that $z^1 \in \mathbb{R}^2_0$ was arbitrary. By setting $z^{1\prime} = z^1$ and solving, we find z^1 and z^2 with the assumed properties and satisfying $z^{1\prime} = z^1$ and $z^{2\prime} = z^2$. This implies that, if we set $W(c^*) = \{z^1, z^2\}$, then $\{z^1, z^2\} \subset \hat{B}^0(c^*, \mathcal{W})$. Recalling that $Y(c^*)$ is the union of all sets of continuation values under external contract c^* associated with normalized self-generating collections, we can thus conclude that $\{z^1, z^2\} \subset Y(c^*)$, and this proves that $\mathrm{Span}(Y(c^*)) = d^*$. Also, because we showed that $d^* \geq \hat{d}$, we thus have established that $\hat{d} = d^*$. $\square$

Step 3: A contractual-equilibrium regime exists, in particular a semi-stationary regime.

Let c^{**} be the semi-stationary external contract that satisfies $g(c^{**}) = \gamma^{**}$ and transitions to c^*. By the argument developed in Subsection 3.2 (and utilized in step 2 above) about the players using ϕ to convexify the set of continuation values, it is the case that

$Y(c^*) = Z(\gamma^*, d^*) + b$ for some $b \in \mathbb{R}^2_0$. This implies that a given action profile α is enforced relative to γ^{**} and y^{**} if and only if α is c^{**}-enforced relative to $\mathcal{Y}$. And then, because γ^{**}, y^{**}, and α^{**} solve Optimization Problem 15 for span d^*, we know that α^{**} solves Optimization Problem 32 under the restriction that the external contract is fixed at c^{**}.

In fact, α^{**} and c^{**} solve Optimization Problem 32 without restriction. To see why the joint value cannot be increased with a different external contract c', note that we can once again utilize the argument developed in Subsection 3.2 to find a semi-stationary external contract c'' that supports the same continuation values and always transitions to c^*. Recall that this conclusion relies on the contractual setting having externally enforced transfers, that only the span of the continuation-value set matters for enforcing action profiles, and that $Y(c^*)$ has the maximal span in collection $\mathcal{Y}$. So in Optimization Problem 32 we can constrain attention to the set of semi-stationary external contracts that transition to c^*. Thus, the only variation is in the stage game specified for the current period, and the maximal joint value is achieved with γ^{**} as in Optimization Problem 15.

Assumption 1 guarantees that $c^{**} \in C$. It is easy to show that Assumption 2 implies $Y(c^0) \neq \emptyset$ by confirming that this set contains $\overline{u}^0(\alpha^0)$, where α^0 is any Nash equilibrium of stage game γ^0. Lemma 5 then establishes the existence of a contractual equilibrium with level $L^* = \overline{u}^{**}_1(\alpha^{**}) + \overline{u}^{**}_2(\alpha^{**})$.

The final step is to construct a semi-stationary contractual equilibrium. From here the steps are similar to those in the proof of Lemma 4, although here we provide details to show that the regime is semi-stationary. The contractual equilibrium we shall construct can be described as a simple three-state system. State 0 refers to histories in H in which the players have not made an agreement to date, so that the inherited external contract was always c^0. State 1 refers to histories with inherited external contract c^* in which the players are coordinating to achieve player 1's least preferred continuation value. State 2 refers to histories with inherited external contract c^* in which the players are coordinating to achieve player 2's least preferred continuation value. The equilibrium continuation values in these three states are denoted, respectively, z^{*0}, z^{*1}, and z^{*2}. We also have the disagreement points for these three states, which are denoted $\underline{w}^0$, $\underline{w}^1$, and $\underline{w}^2$.

These various continuation values are defined as follows. First, we set

$$z^{*1} = \overline{u}^*(\alpha^{*1}) + \frac{\delta}{1-\delta}\overline{y}^{*1}(\alpha^{*1}) + \pi\left(L^* - \overline{u}^*_1(\alpha^{*1}) - \overline{u}^*_2(\alpha^{*1})\right).$$

Note that this implies

$$z^{*1} = (1-\delta)\overline{u}^*(\alpha^{*1}) + \delta\left(z^{*1} + \overline{y}^{*1}(\alpha^{*1})\right) + \pi(1-\delta)\left(L^* - \overline{u}^*_1(\alpha^{*1}) - \overline{u}^*_2(\alpha^{*1})\right).$$

Then we set

$$z^{*2} = (1-\delta)\overline{u}^*(\alpha^{*2}) + \delta\left(z^{*1} + \overline{y}^{*2}(\alpha^{*2})\right) + \pi(1-\delta)\left(L^* - \overline{u}_1^*(\alpha^{*2}) - \overline{u}_2^*(\alpha^{*2})\right).$$

Corresponding to these are disagreement values

$$\underline{w}^1 = (1-\delta)\overline{u}^*(\alpha^{*1}) + \delta\left(z^{*1} + \overline{y}^{*1}(\alpha^{*1})\right)$$

and

$$\underline{w}^2 = (1-\delta)\overline{u}^*(\alpha^{*2}) + \delta\left(z^{*1} + \overline{y}^{*2}(\alpha^{*2})\right).$$

Likewise, we set

$$z^{*0} = \overline{u}^0(\alpha^0) + \pi\left(L^* - \overline{u}_1^0(\alpha^0) - \overline{u}_2^0(\alpha^0)\right),$$

and we note that this implies

$$z^{*0} = (1-\delta)\overline{u}^0(\alpha^0) + \delta z^{*0} + \pi(1-\delta)\left(L^* - \overline{u}_1^0(\alpha^0) - \overline{u}_2^0(\alpha^0)\right).$$

The associated disagreement value is

$$\underline{w}^0 = (1-\delta)\overline{u}^0(\alpha^0) + \delta z^{*0}.$$

It is easy to verify that $z^{*1} = \underline{w}^1 + \pi(L^* - \underline{w}_1^1 - \underline{w}_2^1)$, $z^{*2} = \underline{w}^2 + \pi(L^* - \underline{w}_1^2 - \underline{w}_2^2)$, and $z^{*0} = \underline{w}^0 + \pi(L^* - \underline{w}_1^0 - \underline{w}_2^0)$. To see this, start with each z expression, substitute for $\overline{u}$ using the expression for the disagreement value, and use the fact that $z_1 + z_2 = L^*$. Let us also define

$$w^{**} = (1-\delta)\overline{u}^{**}(\alpha^{**}) + \delta\left(z^{*1} + \overline{y}^{**}(\alpha^{**})\right).$$

This will be the continuation value in from the action phase of a period in which the players selected external contract c^{**} in the negotiation phase. Value w^{**} does not include the transfer made earlier in the current period.

Our semi-stationary contractual equilibrium regime r is specified as follows. For any history h in state 0, where the inherited external contract is c^0, we prescribe $r^{\mathrm{c}}(h) = c^{**}$ and we pick $r^{\mathrm{m}}(h)$ to solve $z^{*0} = r^{\mathrm{m}}(h) + w^{**}$ so that the surplus relative to $\underline{w}^0$ is split according to the bargaining weights. The agreement action profile is $r^{\mathrm{a}}(h, c^{**}, r^{\mathrm{m}}(h)) = \alpha^{**}$, which is enforced by the players using ϕ to randomize between continuation values z^{*1} (moving to state 1) and z^{*2} (moving to state 2) in the following period to achieve the selection $z^{*1} + y^{**}(\cdot)$. For the case of disagreement, we prescribe $r^{\mathrm{a}}(h, c^0, 0) = \alpha^0$, which is enforced by the players coordinating on continuation values z^{*0} again (state 0) in the following period.

For any history h in state 1, where the inherited external contract is c^*, we prescribe $r^{\mathrm{c}}(h) = c^{**}$ and we pick $r^{\mathrm{m}}(h)$ to solve $z^{*1} = r^{\mathrm{m}}(h) + w^{**}$ so that the surplus relative to $\underline{w}^1$ is split according to the bargaining weights. The agreement action profile is $r^{\mathrm{a}}(h, c^{**}, r^{\mathrm{m}}(h)) = \alpha^{**}$, which is enforced just as in the previous paragraph. For the case of disagreement, we prescribe $r^{\mathrm{a}}(h, c^0, 0) = \alpha^{*1}$, which is enforced by the players using ϕ to randomize between continuation values z^{*1} (moving to state 1) and z^{*2} (moving to state 2) in the following period to achieve the selection $z^{*1} + y^{*1}(\cdot)$.

For any history h in state 2, where the inherited external contract is c^*, we prescribe $r^{\mathrm{c}}(h) = c^{**}$ and we pick $r^{\mathrm{m}}(h)$ to solve $z^{*2} = r^{\mathrm{m}}(h) + w^{**}$ so that the surplus relative to $\underline{w}^2$ is split according to the bargaining weights. The agreement action profile is $r^{\mathrm{a}}(h, c^{**}, r^{\mathrm{m}}(h)) = \alpha^{**}$, which is enforced just as in the previous paragraphs. For the case of disagreement, we prescribe $r^{\mathrm{a}}(h, c^0, 0) = \alpha^{*2}$, which is enforced by the players using ϕ to randomize between continuation values z^{*1} (moving to state 1) and z^{*2} (moving to state 2) in the following period to achieve the selection $z^{*1} + y^{*2}(\cdot)$.

By construction, every history in $H(r)$ is in one of the three states described above. Therefore, continuation values are defined for every history in $H(r)$. Regime r is also incentive compatible in the action phase and internally bargain-consistent, again by construction. Furthermore, the regime's level is L^*, which we already found to be the contractual-equilibrium level. Therefore, r is a contractual equilibrium. $\square$

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.

Philippe Aghion, Mathias Dewatripont, and Patrick Rey. Renegotiation design with unverifiable information. *Econometrica*, 62(2):257–282, 1994.

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

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

George Baker, Robert Gibbons, and Kevin J. Murphy. Relational adaptation. Working paper, November 2011. URL http://goo.gl/sci8G.

Daniel Barron, Robert Gibbons, Ricard Gil, and Kevin J. Murphy. Relational adaptation under reel authority. Working paper, 2015.

B. Douglas Bernheim and Michael D. Whinston. Incomplete contracts and strategic ambiguity. *American Economic Review*, 88(4):902–932, 1998.

Jean Beuve and Stéphane Saussier. Interfirm cooperation in strategic relationships: The role of formal contract. *Industrial and Corporate Change*, 21(4):811–836, August 2012.

Patrick Bolton and Mathias Dewatripont. *Contract Theory*. MIT, Cambridge, M.A., 2005.

C. Bull. The existence of self-enforcing implicit contracts. *Quarterly Journal of Economics*, 109 (1):147–159, 1987.

Kristy Buzard and Joel Watson. Contract, renegotiation, and hold up: Results on the technology of trade and investment. *Theoretical Economics*, 7:283–322, 2012.

Yeon-Koo Che and Donald B. Hausch. Cooperative investments and the value of contracting. *American Economic Review*, 89(1):125–147, 1999.

Yeon-Koo Che and S.-W. Yoo. Optimal incentives for teams. *American Economic Review*, 91(3): 525–541, 2001.

Aaron S. Edlin and Stefan Reichelstein. Holdups, standard breach remedies, and optimal investment. *American Economic Review*, 86(3):478–501, 1996.

R. Evans. Simple efficient contracts in complex environments. *Econometrica*, 76(3):459–491, 2008.

Yuk-fai Fong and Jin Li. Relational contracts, limited liability, and employment dynamics. *Journal of Economic Theory*, 169:270–293, 2017.

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

Oliver Hart and John Moore. Incomplete contracts and renegotiation. *Econometrica*, 56(4):755–785, 1988.

Oliver D. Hart and John Moore. Foundations of incomplete contracts. *Review of Economic Studies*, 66:115–138, 1999.

Elisabetta Iossa and Giancarlo Spagnolo. Contracts as threats: On a rationale for rewarding A while hoping for B. Working paper, 2011. URL http://goo.gl/Fut6K.

Hideshi Itoh and Hodaka Morita. Formal contracts, relational contracts, and the threat-point effect. *American Economic Journal: Microeconomics*, 7(3):318–46, August 2015.

B. Klein and K.B. Leffler. The role of market forces in assuring contractual performance. *Journal of Political Economy*, 89(4):615–641, 1981.

Rumen Kostadinov. Renegotiation of long-term contracts with implicit incentives. Working paper, 2017.

Ola Kvaløy and Trond E. Olsen. Endogenous verifiability and relational contracting. *American Economic Review*, 99(5):2193–2208, 2009.

Sergio G. Lazzarini, Gary J. Miller, and Todd R. Zenger. Order with some law: Complementarity versus substitution of formal and informal arrangements. *Journal of Law, Economics, and Organization*, 20(2):261–298, October 2004.

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

S. Macaulay. Non-contractual relations in business: A preliminary study. *American Sociological Review*, 28:55–67, 1963.

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

I.R. Macneil. Contracts: Adjustment of long-term economic relations under classical, neoclassical, and relational contract law. *Northwestern University Law Review*, 72(6):854–904, 1978.

James M. Malcomson. Relational incentive contracts. In Robert Gibbons and John Roberts, editors, *The Handbook of Organizational Economics*, chapter 25, pages 1014–1065. Princeton University Press, 2013.

David Martimort, Aggey Semenov, and Lars Stole. A theory of contracts with limited enforcement. *Review of Economic Studies*, 2016.

Eric Maskin and Jean Tirole. Unforeseen contingencies and incomplete contracts. *Review of Economic Studies*, 66:83–113, 1999.

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

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

Georg Noldeke and Klaus M. Schmidt. Option contracts and renegotiaton: A solution to the holdup problem. *Rand Journal of Economics*, 26:163–179, 1995.

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

Laura Poppo and Todd Zenger. Do formal contracts and relational governance function as substitutes or complements? *Strategic Management Journal*, 23(8):707–725, Aug. 2002.

Debraj Ray. The time structure of self-enforcing agreements. *Econometrica*, 70(2):547–582, March 2002.

Michael D. Ryall and Rachelle C. Sampson. Formal contracts in the presence of relational enforcement mechanisms: Evidence from technology development projects. *Management Science*, 55 (6):906–925, June 2009.

Mikhail Safronov and Bruno Strulovici. Strategic renegotiation in repeated games. Working paper, 2016.

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

Ilya Segal. Complexity and renegotiation: A foundation for incomplete contracts. *Review of Economic Studies*, 66:57–82, 1999.

Carl Shapiro and Joseph E. Stiglitz. Equilibrium unemployment as a worker discipline device. *The American Economic Review*, 74(3):433–444, 1984.

Joel Watson. Contract, mechanism design, and technological detail. *Econometrica*, 75(1):55–81, 2007. ISSN 00129682, 14680262. URL http://www.jstor.org/stable/4123108.

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

Relational Contracting, Negotiation, and External Enforcement

David Miller, Trond E. Olsen, and Joel Watson

March 2018

B Supplementary Appendix

This supplemental appendix contains an additional existence theorem and proof, notes on the result that strengthening external enforcement leads to a weakly higher contractual equilibrium level, and detailed analysis for some of the examples in Section 4.

B.1 Existence in Finite Settings

In this subsection, we provide an existence result for settings with finite stage games and a finite set C. Here the other aspects of the relational contracting games are fully general, so we do not need to make Assumptions 1–3 and we are not constraining attention to settings with externally enforced transfers.

Theorem 4. *For any relational-contract setting in which C is finite and every game in G is finite, a contractual equilibrium exists.*

Proof of Theorem 4. We start by proving that there is a collection of continuation-value sets $\mathcal{W}^0$ that satisfies $W^0(c) \subset B^0(c, \mathcal{W}^0)$ for all $c \in C$. In particular, we will work with collections that have singleton continuation-value sets: for each $c \in C$, $W^0(c) = \{w^c\}$, where $w^c \in \mathbb{R}^2$. Note that $w_1^c + w_2^c = 0$ for all $c \in C$, so we can think of these points as being on the real line.

For any point $\nu = (w^c)_{c \in C} \in \mathbb{R}^{|C|}$ that defines $\mathcal{W}^0$ by $W^0(c) = \{w^c\}$ for all $c \in C$, let $f(\nu) \equiv \prod_{c \in c} \operatorname{Conv} B^0(c, \mathcal{W}^0)$, where "Conv" denotes the convex hull. Because the stage games are finite, the bargaining solution maps supported values to the zero-value line along the ray π, and because continuation values are discounted, we can find a bound κ such that $w^c \in [-\kappa, \kappa]^2$ for all $c \in C$ implies that $B^0(c, \mathcal{W}^0) \subset [-\kappa, \kappa]^2$. Further, because each stage game is finite and the Nash correspondence is nonempty and upper hemi-continuous in payoff vectors, B^0 has the same property. Thus, f is a correspondence from a compact set to itself, it is nonempty and convex valued, and it is upper-hemicontinuous. By the Kakunati fixed-point theorem, f has a fixed point $\overline{\nu} = (\overline{w}^c)_{c \in C}$.

Let $\overline{\mathcal{W}} = \{\overline{W}(c)\}_{c \in C}$ be defined by $\overline{W}(c) = \{\overline{w}^c\}$ for all $c \in C$. The fixed point property means that $\overline{W}(c) \subset \operatorname{Conv} B^0(c, \overline{\mathcal{W}})$ for all $c \in C$, but it is not necessarily the case that $\overline{W}(c) \subset B^0(c, \overline{\mathcal{W}})$ for all $c \in C$. However, if this latter condition fails, then we can find two points $\overline{w}'^c, \overline{w}''^c \in B^0(c, \overline{\mathcal{W}})$ such that $\overline{w}^c$ is on the line between $\overline{w}'^c$ and $\overline{w}''^c$. The players can achieve an expected continuation value of $\overline{w}^c$ in a period in which the inherited external contract is c, from the perspective of the previous period by using the random draw ϕ to randomize between $\overline{w}'^c$ and $\overline{w}''^c$.[30] For each affected external contract $c \in C$, we replace $\overline{W}(c) = \{\overline{w}^c\}$ with $\overline{W}(c) = \{\overline{w}'^c, \overline{w}''^c\}$, and the adjusted $\overline{\mathcal{W}}$ satisfies $\overline{W}(c) \subset B^0(c, \overline{\mathcal{W}})$ for all $c \in C$. Because $\overline{\mathcal{W}}$ is "fully nonempty", meaning that $W(c) \neq \emptyset$ for all c, continuation values from the action phase are supported (for every external contract) and there is a level L such that $\overline{\mathcal{W}} + L\pi$ is self-generating with level L. This implies that $\mathcal{Y}$ is also fully nonempty.

To complete the proof, we must show that Optimization Problem 32 has a solution. By upper hemi-continuity of B^0 and that $\mathcal{Y}$ is normalized self-generating, we know that the closure of $\mathcal{Y}$, denoted by $\operatorname{Clos} \mathcal{Y}$, is also normalized self-generating. Here, $\operatorname{Clos} \mathcal{Y} = \{\operatorname{Clos} Y(c)\}_{c \in C}$. This means that $\mathcal{Y} = \operatorname{Clos} \mathcal{Y}$. Thus, for each $c \in C$, the problem of maximizing $u_1(\alpha; c) + u_2(\alpha; c)$ over all c-enforced action profiles $\alpha \in \Delta A(c)$ has a solution. Because there are a finite number of external contracts, the overall maximum exists. $\square$

B.2 Stronger enforcement technologies

In this subsection we describe how to extend Theorem 1 to compare external enforcement technologies without requiring inclusion. Here is a weaker definition of "stronger external enforcement technology" than is described in the text:

Definition 13. *External contracting environment $\mathcal{A} = (G', C', g', c^{0\prime}, \zeta')$ is* ***directly stronger*** *than external contracting environment $\mathcal{B} = (G, C, g, c^0, \zeta)$ if any of the following are true:*

- *Relabel: $(G', C', g', c^{0\prime}, \zeta')$ is isomorphic to (G, C, g, c^0, ζ);*
- *Enlarge: $G' \supset G$, $C' \supset C$, $g'(c) = g(c)$ for all $c \in C$, $c^{0\prime} = c^0$, and $\zeta'(c, x, \phi) = \zeta(c, x, \phi)$ for all $c \in C$, $x \in X(c)$, and $\phi \in [0, 1]$;*
- *Refine: There is a one-to-one mapping from G to G' such that for each $(A, X, \lambda, u, P) \in G$, there exists $(A, X, \lambda, u, P') \in G'$ such that P' is finer than P; and, letting stage games mapped to each other have the same labels, $g'(c) = g(c)$ for all $c \in C$, $c^{0\prime} = c^0$, and $\zeta'(c, x, \phi) = \zeta(c, x, \phi)$ for all $c \in C$, $x \in X(c)$, and $\phi \in [0, 1]$;*

[30] Because ϕ is uniformly distributed, for any outcome that would lead to external contract c for a positive mass of the random draw, the players can divide this set of ϕ values to achieve any probability distribution over the continuation values $\overline{w}'^c$ and $\overline{w}''^c$.

*$\mathcal{A}$ is **stronger** than $\mathcal{B}$ if there exists a sequence $\{\mathcal{C}\}_{k=1}^{K}$ such that $\mathcal{A}$ is directly stronger than $\mathcal{C}_1$, $\mathcal{C}_k$ is directly stronger than $\mathcal{C}_{k+1}$ for all $k = 1, \ldots, K-1$, and $\mathcal{C}_K$ is directly stronger than $\mathcal{B}$.*

Proof of Theorem 1 with "stronger" defined here. Relabeling via an isomorphism clearly has no effect on the available equilibria.

Enlarging the range of enforcement options preserves the incentive compatible and internal bargain-consistency of regimes. Therefore the supremal level among incentive compatible and internally bargain-consistent regimes must weakly increase.

Refining enforcement capabilities by making the enforcer's partition finer takes more care to analyze, since it changes the set of histories $H(r)$ that are available under a regime. Let r be a contractual equilibrium in the environment described in Section 2.1, which we will refer to as the environment with partition P. Consider an alternative environment where every stage game (A, X, λ, P) in G is replaced by (A, X, λ, P'), where P' is weakly finer than P. Let (G', C', g', ζ') be the elements corresponding to (G, C, g, ζ) in this environment, (i.e. the set of stage games, the set of external contracts, the mapping $g' : C' \to G'$, and the transition function, respectively). The transition function ζ' must now be measurable with respect to P' in the same sense as ζ is measurable with respect to P.

First note that a contract $c \in C$ is feasible under P', and thus an element of C' in the following sense. If $g(c)$ is the game (A, X, λ, P) under P, define $g'(c)$ to be the game (A, X, λ, P') under P'. The transition ζ' for contract c is then measurable with respect to P'. This follows because for any two outcomes $x, x' \in X(c)$ with $x' \in P'(x, c)$, and therefore $x' \in P(x, c)$, we have $\zeta(c, x', \phi) = \zeta(c, x, \phi)$ by P-measurability of ζ. Thus, any two outcomes in the same partition element of P' transition to the same continuation contract, and the transition $\zeta'(c, \cdot, \phi)$ is therefore measurable with respect to $P'(\cdot; c)$.

The regime r is defined on histories $h \in H$ under P, where a T-period history is a sequence $h = \left\{(c^t, m^t, x^t, \phi^t)\right\}_{t=1}^{T}$ as defined in Section 2.3, with $c^t \in C$. Given that a contract in C is also in C' (in the sense just explained), regime r is also well defined for histories $h' \in H'$ under P' that contain only contracts from C, i.e., histories of the form $\left\{(c^t, m^t, x^t, \phi^t)\right\}_{t=1}^{T}$ with $c^t \in C$. Moreover, for any such T-period history, regime r will in period $T+1$ select a contract $c^{T+1} \in C$ if there is agreement, and continue with $c^T \in C$ if there is disagreement. The set $H(r)$ of histories in which, in each period, either the players made the agreement specified by the regime r or there was disagreement, will therefore remain the same under P' as under P. Then we see from Definitions 1 and 2 that regime r will satisfy incentive compatibility and internal bargaining consistency under the finer partition P'. Moreover, this will be true irrespective of how we extend the definition of regime r to histories $h' \in H'$ that contain contracts in C' but not in C. $\square$

B.3 Monitoring and options

Manager control Consider the case where the manager has the right to select from the options. For a given option contract the analysis proceeds as in Section 1.2. Consider first the worst disagreement value for the worker (player 1). Let the regime here call for the manager to select (p^1, μ^1), the worker to exert effort, and for the parties to coordinate on continuation value $z^1 + (\rho, -\rho)$ if the monitor signal is high and (p^1, μ^1) was selected, and on z^1 otherwise. Given that (p^1, μ^1) is selected, the worker will then exert effort as intended when

$$\delta\rho \geq (1-\delta)\beta/\mu^1,$$

and the disagreement value that is worst for player 1 will be given by

$$\underline{v}^1 = (1-\delta)(-\beta - p^1, 1 - k(\mu^1) + p^1) + \delta z^1 + \delta(\rho, -\rho), \tag{36}$$

with ρ minimal, thus $\rho = \frac{1-\delta}{\delta}\frac{\beta}{\mu^1}$. Incentive constraints for the manager's selection will be considered below.

Let $L^1 = 1 - \beta - k(\mu^1)$ be the welfare generated in the disagreement period. The parties will negotiate to avoid disagreement, and the equilibrium payoff that is worst for the worker will be $z^1 = \underline{v}^1 + \pi(L - \underline{v}^1_1 - \underline{v}^1_2)$. A little algebra yields

$$z^1 = (-p^1 + \beta/\mu^1 - \beta, 1 - k(\mu^1) + p^1 - \beta/\mu^1) + \pi(L - L^1) \tag{37}$$

The term $\beta/\mu^1 - \beta$ is the rent accruing to the worker from his effort under imperfect monitoring.

Consider next the disagreement point that is worst for player 2. Here the regime calls for the worker to shirk, the manager to select the option (p^2, μ^2), and for the parties to coordinate on z^2 for any outcome. The manager is then willing to select the appropriate option provided $p^2 - k(\mu^2) \geq p^1 - k(\mu^1)$. This yields disagreement value

$$\underline{v}^2 = (1-\delta)(-p^2, p^2 - k(\mu^2)) + \delta z^2 \tag{38}$$

In equilibrium negotiations will prevent disagreement and lead to payoffs $z^2 = \underline{v}^2 + \pi(L - \underline{v}^2_1 - \underline{v}^2_2)$, which can now be written as

$$z^2 = (-p^2, p^2 - k(\mu^2)) + \pi(L - L^2). \tag{39}$$

Here $L^2 = -k(\mu^2)$ is the one-period welfare level should such a disagreement occur. This verifies the payoff expressions given in (16) in the text.

It follows that the span $d = z_1^2 - z_1^1$ is given by

$$d = p^1 - p^2 - (\beta/\mu^1 - \beta) + \pi_1(L^1 - L^2) \tag{40}$$

where $L^1 - L^2 = 1 - \beta - k(\mu^1) + k(\mu^2)$, and IC for the manager's selection of (p^2, μ^2) requires $k(\mu^1) - k(\mu^2) \geq p^1 - p^2$. We see that the span is maximal when $p^1 - p^2$ is maximal (and thus when (17) holds)—which verifies (18) in the text— and consequently when μ^1 is maximal and μ^2 minimal ($\mu^1 = 1$, $\mu^2 = 0$). The maximal span is thus

$$d = (k(1) - k(0))(1 - \pi_1) + \pi_1(1 - \beta)$$

The options contract results in a larger span than what is obtained without such a contract, where the monitor level is inefficiently high in all cases under disagreement; and the span is $\pi_1(1-\beta)$. The latter situation implies a larger welfare difference $L^1 - L^2$, which in isolation yields a larger span, but this is more than compensated for in the options contract via the payment difference $p^1 - p^2$.

It remains to verify that the manager has no incentives to deviate from selecting the option (p^1, μ^1) when the worker is to be punished. Note that compliance gives payoff $\underline{v}_2^1 = (1-\delta)(1-k(\mu^1)+p^1-\beta/\mu^1)+\delta z_2^1$. A deviation to (p^2, μ^2) with $\mu^2 = 0$ would make the worker shirk, and the manager's payoff would then be $(1-\delta)(-k(\mu^2)+p^2)+\delta z_2^1$. The latter is smaller than $\underline{v}_2^1$ due to (17) and $\beta/\mu^1 = \beta < 1$. This verifies that all IC constraints are satisfied.[31]

Worker control Now consider the case where the worker has the right to select from an options contract.

Consider first the disagreement point that is worst for the manger (player 2). The regime here calls for the worker to select the option (p^2, μ^2), then to shirk, and the parties to coordinate on z^2 for any outcome. These actions for the worker are incentive compatible if $p^2 \leq p^1$, and the payoffs are then as above given by (38) and (39) under disagreement and agreement, respectively.

Next consider the disagreement point that is worst for player 1, where it is intended that the worker selects (p^1, μ^1) and exerts effort $a = 1$. Let the regime here call for coordination

[31] It is straightforward to verify that participation constraints for the worker ($\underline{v}_1^1, \underline{v}_1^2 \geq 0$) can be satisfied by e.g. setting p^1 such that $\underline{v}_1^1 = 0$ and $p^2 = p^1 - k(1) + k(0)$.

on $z^1 + (\rho, -\rho)$ if the worker selects (p^1, μ^1) and the signal is high, and on z^1 otherwise, where as above $\rho = \frac{1-\delta}{\delta}\frac{\beta}{\mu^1} \leq d$. For the given option, effort is then incentive compatible and leads as above to disagreement values given by (36).

If the worker deviates and selects the other option (p^2, μ^2), he will optimally shirk, and thus get payoff $(1-\delta)(-p^2) + \delta z_1^1$. Option (p^1, μ^1) is thus his best choice if $\delta\rho \geq (1-\delta)(p^1 - \beta - p^2)$, i.e. if $\beta(1/\mu^1 - 1) \geq p^1 - p^2$.

Negotiations then yield values z^1 as in (37), and consequently a span given by (40) above. But now the IC constraints for the worker's selection of options are $\beta(1/\mu^1 - 1) \geq p^1 - p^2 \geq 0$, and we must moreover have $\rho = \frac{1-\delta}{\delta}\frac{\beta}{\mu^1} \leq d$. From (40) we see that the largest span is obtained when $p^1 - p^2$ is maximal, and thus here $d = \pi_1(L^1 - L^2) = \pi_1(1 - \beta - k(\mu^1) + k(\mu^2))$. This verifies (19) and (20) in the text.

We also note that all of this can be implemented without violating participation constraints for the worker by e.g. setting $p^2 = 0$ and $p^1 = \beta\frac{1-\mu^1}{\mu^1}$.

Under worker control, incentive constraints for the selection of options imply that the span is proportional to the welfare difference $L^1 - L^2$; and thus largest when the monitor level μ^2 is maximal. This is highly inefficient, since (absent agreement) μ^2 will be implemented when no effort is to be provided. And it is quite the opposite of the case of manager control considered above, where the largest feasible span is obtained with inefficiently high monitoring when the worker is supposed to exert effort under disagreement. The differences between the two cases reflect the differences in the two parties' incentives when choosing between options.

B.4 Multitasking

Worker control Consider first the case where the worker has the right to select among the options (a_2^i, p^i), $i = 1, 2$. Under disagreement to punish the manager, let the regime call for the worker to supply $a = (0, a_2^2)$, and the parties to coordinate on z^2 next period for every outcome this period. This is incentive compatible for the worker as long as $p^2 - \kappa(0, a_2^2) \geq p^1 - \kappa(0, a_2^1)$, and leads to disagreement values

$$\underline{v}^2 = (1-\delta)(p^2 - \kappa(0, a_2^2), \nu(0, a_2^2) - p^2) + \delta z^2 \tag{41}$$

Under disagreement to punish the worker, let the regime call for the worker to supply $a = (a_{1h}, a_2^1)$, and the parties to coordinate on $z^1 + (\rho, -\rho)$ unless the worker deviates, in which case they coordinate on z^1. For given quantity a_2^1, high quality is incentive compatible for the worker if

$$(1-\delta)(p^1-\kappa(a_{1h},a_2^1))+\delta(z_1^1+\rho)\geq(1-\delta)(p^1-\kappa(0,a_2^1))+\delta z_1^1$$

with $\rho \leq d = z_1^2 - z_1^1$. To maximally punish the worker, ρ should be minimal and thus given by

$$(\kappa(a_{1h},a_2^1)-\kappa(0,a_2^1))(1-\delta)=\delta\rho \tag{42}$$

Selecting quantity a_2^1 is then incentive compatible for the worker if $p^1-\kappa(0,a_2^1)\geq p^2-\kappa(0,a_2^2)$. This leads to disagreement values

$$\underline{v}^1=(1-\delta)(p^1-\kappa(0,a_2^1)),\nu(a_{1h},a_2^1)-\kappa(a_{1h},a_2^1)-(p^1-\kappa(0,a_2^1)))+\delta z^1 \tag{43}$$

Negotiations yield $z^i=\underline{v}^i+\pi(L-\underline{v}_1^i-\underline{v}_2^i)$, and thus

$$\begin{aligned}
z_1^2 &= (p^2-\kappa(0,a_2^2))+\pi_1(L-(\nu(0,a_2^2)-\kappa(0,a_2^2))\\
z_1^1 &= (p^1-\kappa(0,a_2^1))+\pi_1(L-(\nu(a_{1h},a_2^1)-\kappa(a_{1h},a_2^1))
\end{aligned}$$

Recall that incentive compatibility requires $p^1-\kappa(0,a_2^1)\geq p^2-\kappa(0,a_2^2)$, hence we see that the span $d=z_1^2-z_1^1$ is largest when this constraint binds and $a_2^2=0$. The span is then

$$d=\pi_1(\nu(a_{1h},a_2^1)-\kappa(a_{1h},a_2^1))$$

This implies $p^1-p^2=\kappa(0,a_2^1)$, and we can set $p^2=\kappa(0,0)=0$. Thus, the court enforced payment compensates the worker for the cost of providing low quality of the selected quantity. Moreover, the optimal a_2^1 is the maximal quantity (of the high quality good) that can be implemented with the equilibrium span, thus it coincides with the equilibrium quantity a_2 under agreement. This verifies the assertions in the text regarding worker control of options.

Manager control Assume next that the manager has the right to select among the options. Under disagreement to punish the manager (reward the worker), let the regime then call for the manager to select option (a_2^2, p^2), the worker to provide $a=(0,a_2^2)$, and for the parties continue with z^2 next period for any outcome this period. This is incentive compatible provided $\nu(0,a_2^2)-p^2\geq\nu(0,a_2^1)-p^1$ with $p^2\geq\kappa(0,a_2^2)$, and leads to disagreement values as in (41) above.

Under disagreement to punish the worker, let the regime call for the manager to select option (a_2^1, p^1) and the worker to provide $a = (a_{1h}, a_2^1)$. Let the regime also call for coordination on z^1 if only the worker deviates, and on $z^1 + (\rho, -\rho)$ otherwise, where $\rho \geq 0$ is given by (42).

Given the option (a_2^1, p^1), high quality is then incentive compatible for the worker. We must of course have $\rho \leq d$, and thus $a_2^1 \leq a_2$ (the quantity supplied under agreement). Note that by setting $\rho = 0$, we may allow $a_2^1 = 0$.

The manager's choice of option is incentive compatible if $\nu(a_{1h}, a_2^1) - p^1 \geq \nu(0, a_2^2) - p^2$. (If she deviates, the worker will supply low quality, and the manager will then be worse off.) Substituting for ρ, we then see that disagreement values are here given as in (43) above.

By internal bargaining consistency this leads to the same expressions as above for the values z_1^1, z_1^2, but incentive compatibility for the manager now requires

$$\nu(a_{1h}, a_2^1) - p^1 \geq \nu(0, a_2^2) - p^2 \geq \nu(0, a_2^1) - p^1$$

The span $d = z_1^2 - z_1^1$ is then largest for payments such that $p^2 - p^1 = \nu(0, a_2^2) - \nu(0, a_2^1)$, which yields

$$\begin{aligned} d &= (1 - \pi_1)(\nu(0, a_2^2) - \kappa(0, a_2^2)) \\ &\quad + \pi_1(\nu(a_{1h}, a_2^1) - \kappa(a_{1h}, a_2^1)) - (\nu(0, a_2^1) - \kappa(0, a_2^1)) \end{aligned}$$

The maximal span with these options is obtained by choosing $a_2^2 = \arg\max_{a_2}(\nu(0, a_2) - \kappa(0, a_2)) \equiv a_2^0$, and a_2^1 to maximize the expression in the last line, hence we have

$$\begin{aligned} d &= d(a_2) = (1 - \pi_1)\left(\nu(0, a_2^0) - \kappa(0, a_2^0)\right) \\ &\quad + \max_{0 \leq a_2' \leq a_2} \left[\pi_A(\nu(a_{1h}, a_2') - \kappa(a_{1h}, a_2')) - (\nu(0, a_2') - \kappa(0, a_2'))\right] \end{aligned}$$

The incentive compatible payments can be set as $p^i = \nu(0, a_2^i)$, $i = 1, 2$, implying that the worker is paid the gross value of low quality in this scheme.

The equilibrium quantity of the (high quality) good is now given by the largest solution to [32]

$$\kappa(a_{1h}, a_2) - \kappa(0, a_2) = \frac{\delta}{1 - \delta} d(a_2).$$

Clearly the equilibrium span here is larger than the corresponding span under worker control

[32]It may be noted that, if the maximal span $d(a_2)$ is obtained for $a_2' = a_2$, then the externally enforced option contract need not be renegotiated in any period, since it will implement high-quality quantity a_2 also in agreement. Otherwise it will be renegotiated under agreement.

when π_1 is sufficiently small. This verifies the assertions in the text.

B.5 Partnership

In the partnership game, whenever efficiency is attainable the contractual equilibrium maximizes the span. The solution is semistationary by Theorem 3, with the external contract specifying $T^* = 0$ in the current period followed by $\hat{T} \in [0, 1]$ in all future periods. We show here that in contractual equilibrium $\hat{T} = 1$ if efficiency is attainable.

Consider a history off the equilibrium path, when partner 1 is supposed to be punished and the players have just disagreed. Their regime calls for play of $a = (1, 0)$ in the stage game with $\hat{T}$; this is enforced by continuing with value $z^1 + (\rho, -\rho)$ if $a \in \{(1, 0), (1, 1), (0, 1)\}$ is played, and z^1 if $a = (0, 0)$ is played. The incentive compatibility constraint for partner 1 is

$$\underline{v}_1^1 = (1 - \delta)(1 - T)(-\sigma) + \delta(z_1^1 + \rho) \geq 0 + \delta z_1^1. \tag{44}$$

The incentive compatibility constraint for partner 2 is satisfied automatically, since $a_2 = 0$ is a best response to $a_1 = 1$ in the stage game for all $T \in]0, 1]$; the continuation value for player 2 is

$$\underline{v}_2^1 = (1 - \delta)\big((1 - T) + T(\xi - \beta)\big) + \delta(L - z_1^1 - \rho). \tag{45}$$

Now step back to the start of the period. Knowing that $\underline{v}^1$ is what they will get if they disagree, they renegotiate to the payoff vector z^1, characterized by

$$z^1 = \underline{v}^1 + \frac{1}{2}\big(L - \underline{v}_1^1 - \underline{v}_2^1, L - \underline{v}_1^1 - \underline{v}_2^1\big). \tag{46}$$

Moreover, $z^1 + (\rho, -\rho)$ must be contained in the equilibrium value set:

$$z_1^1 + \rho \leq L - z_1^1 \tag{47}$$

For any δ, if efficiency is attainable then the contractual equilibrium selects ρ and $\hat{T}$ to minimize z_1^1 subject to these constraints, along with the equilibrium-path incentive constraint:

$$1 - \beta \geq (1 - \delta) + \delta z_1^1. \tag{48}$$

For fixed δ and T, a necessary condition for minimizing z_1^1 subject to these constraints is to choose ρ to bind Eq. (44), which yields

$$\rho = \frac{(1-\delta)(1-T)}{\delta}\sigma. \tag{49}$$

Then

$$z_1^1 = \frac{1 - 2\beta + \sigma + T(1 + \beta - \xi - \sigma)}{2}, \tag{50}$$

which is decreasing in T under these assumptions. Choosing $T = 1$ minimizes z_1^1 subject to Eq. (44), and satisfies both Eq. (47) and Eq. (48) if $\delta \geq \frac{2\beta}{\beta+\xi} < 1$.