American Economic Review 2020, 110(7): 2153–2197
https://doi.org/10.1257/aer.20180427

Relational Contracting, Negotiation, and External Enforcement†

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

We study relational contracting and renegotiation in environments with external enforcement of long-term contractual arrangements. A long-term contract governs the stage games that 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, jointly optimize when selecting a contract, and exercise their relative bargaining power. Our main result is that in a wide variety of settings, the optimal contract is* ***semi-stationary****, with stationary terms for all future periods but special terms for the current period. In each period the parties renegotiate to this same contract. 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.* (*JEL* C73, C78, D23, D86)

In many long-term relationships, such as between a worker and a firm, two business partners, or an upstream supplier and a downstream buyer, the parties would like to cooperate for their mutual benefit but are each tempted to deviate for individual gain. The contracts they form typically provide incentives through a combination of *self-enforcement* (the parties' coordinated behavior to reward and punish each other over time) and *external enforcement*, such as provided by courts and the legal system. The literature on relational contracting has provided insights on self-enforcement in the context of stationary externally enforced terms. We develop a general model in which the parties can write arbitrary nonstationary, long-term

* Watson: UC San Diego (email: jwatson@ucsd.edu); Miller: University of Michigan (email: econdm@umich.edu); Olsen: NHH Norwegian School of Economics (email: trond.olsen@nhh.no). Jeffrey Ely was the coeditor for this article. We depart from alphabetical order to acknowledge the first author's outsized conceptual and technical contributions to this project. For helpful comments and encouragement, the authors thank Jeff Ely and five anonymous referees; Nageeb Ali, Daniel Barron, Matthias Fahn, and Andy Skrzypacz; organizers and participants at the 2016 Econometric Society Winter Meetings, the 2016 Stanford Institute for Theoretical Economics summer conference, the Third Annual Conference on Relational Contracts at Northwestern (2017), the 2018 American Economic Association Winter Meetings, the 2018 EARIE Conference, and the Contract Theory Workshop at Kyoto (2019); as well as seminar participants at Cambridge INET Institute, Florida International University, Indiana University, Johns Hopkins Carey Business School, Kellogg, Konstanz, Michigan, NHH Norwegian School of Economics, Penn State, Princeton, Seoul National University, Shanghai Jiao Tong University, Tulane, UCSD, UC3 Madrid, and Washington. Watson thanks the NSF (SES-1227527) and UCSD for grant support.
† Go to https://doi.org/10.1257/aer.20180427 to visit the article page for additional materials and author disclosure statements.

contracts that they can freely renegotiate at any time. We provide results on the form of optimal contracts and on the complementarity of external enforcement and self-enforcement. Further, we present novel applications in which a worker and manager contract over a monitoring technology.

The prior literature establishes that, in a stationary environment without external enforcement, if the parties can make monetary transfers that enter their payoffs linearly, then stationary behavior on the equilibrium path is optimal (see, e.g., Levin 2003, Miller and Watson 2013). Introducing external enforcement, we find that while the parties optimally write the same long-term contract every time they renegotiate, the contract they write is in general nonstationary. If monetary transfers as a function of verifiable outcomes can be externally enforced, or if no outcomes are verifiable, then the nonstationarity takes a particular form with regard to external enforcement: the future part of the contract, which the parties will inherit in the next period, is stationary; but the present part, which governs the current period, is special. We call such a contract *semi-stationary*. Intuitively, the parties choose the future part to maximize the power of incentives, while they choose the present part to maximize their joint payoffs given the power of incentives available to them. Since they anticipate renegotiating in each new period, along the equilibrium path they always operate under the present part of the optimal contract.[1]

A common theme in our applications is that, because equilibrium contracts are semi-stationary, strict contractual terms such as intense monitoring are routinely adjusted to milder terms in the short run. Such behavior is often observed in reality. For instance, many organizations have strict formal rules, regarding attendance and procedures at work, that management routinely allows employees to bend. Our result on complementarity speaks to empirical findings as well.[2]

Following the relational contracts literature (e.g., Levin 2003, Malcomson 2013), we view the *contract* between parties as an agreement encompassing both externally enforced and self-enforced parts. The former, which we call the *external contract*, prescribes how a court or other external referee is to intervene in the relationship conditional on verifiable information. The latter, self-enforced part specifies the parties' individual actions over time, as well as their anticipated revisions of the external contract.[3] In our model, the external contract specifies the stage game to be played in each period, as a function of the verifiable history. We normally refer to an external contract as simply a *contract*, as it will typically be clear from the context whether we are addressing both parts of the contract or just the external part. We add "external" where needed to avoid confusion.

Allowing for arbitrary long-term contracts sets our model apart from the previous literature on relational contracting with limited external enforcement, which has typically either allowed for only short-term (spot) contracts, or assumed that long-term

[1] While a semi-stationary contract is intended to be renegotiated every period, we explain in Section IIIA how such renegotiation could be avoided in an expanded model in which contracts can include options, without otherwise affecting any of our conclusions.

[2] Iossa and Spagnolo (2011) provides an explanation of the first phenomenon that is related to ours; we discuss the differences in Section V. Empirical findings of complementarity are briefly discussed in Section IIID.

[3] See Watson (2013, 2001) for detailed discussion. In the literature, external and self-enforced contractual elements are variously differentiated with the terms "explicit/implicit," "formal/informal," and "legal/relational." The "external/self-enforced" terminology we prefer focuses attention on the source of the enforcement power.

contracts are stationary.[4] Though the environment is stationary, nonstationary contracts introduce a payoff-relevant state variable: the contract the parties agreed upon previously sets the disagreement point for renegotiation in the current period, and it therefore influences the agreement that they reach. Moreover, much of the related literature assumes that self-enforcement is irrevocably terminated after a deviation, so then parties behave myopically. In contrast, we suppose that the parties can renegotiate all aspects of their relationship every period, and we find that they continue to combine optimal self-enforcement with external enforcement even after a deviation. Our approach thus addresses how agents initiate and manage their relationship, including how their agreements evolve after deviations and disagreements.

Kostadinov (2019), developed independently and contemporaneously, is the only other project of which we are aware that studies long-term, nonstationary contracts in an environment with external enforcement and renegotiation. That project is conceptually distinct from ours on two dimensions. First, Kostadinov (2019) studies a particular principal-agent game in which the agent is strictly risk averse, and uses the specific properties of the agent's risk aversion to prove results. In contrast, we examine a wide range of settings with monetary transfers and quasilinear utility. Second, Kostadinov (2019) primarily employs a concept of renegotiation based on "strong optimality" (following Levin 2003), without a theory of bargaining. In contrast, we use *contractual equilibrium* to explicitly model renegotiation, as we describe next. Nonetheless, Kostadinov (2019) finds a comparable result: in a strongly optimal equilibrium, the long-term contract is renegotiated each period and is nonstationary.[5]

Our model applies the concept of *contractual equilibrium* (Miller and Watson 2013, Watson 2013) to a hybrid dynamic game in which each period has first a cooperative *negotiation phase* and then a noncooperative *action phase*. In the negotiation phase, players renegotiate their contract and can make monetary transfers; in equilibrium they reach an agreement according to the generalized Nash (1950) bargaining solution. The disagreement point entails no immediate transfer. In the action phase, players choose actions in the contractually specified stage game; in equilibrium these actions depend only on the public history and satisfy individual incentive constraints, as in a perfect public equilibrium.[6]

Our model accommodates a variety of applications (such as employment relations, repeated procurement, team production, and partnerships) and a variety of externally enforced elements, such as contingent payments, production technologies, and task assignment. In an application, the scope of external enforcement is represented by an exogenously given set of stage games that are available for the players to specify in their contract. Each stage game includes a partition defining the extent to which outcomes are verifiable.

[4] Prominent entries include 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).

[5] The Kostadinov (2019) logic is similar to that behind our main result: players design the future part of their contract to harshly punish a deviating player, but each period they renegotiate to special terms for the current period. Kostadinov (2019) also shows that such a result would also arise from applying a generalized notion of contractual equilibrium in his model.

[6] Miller and Watson (2013) and Watson (2013) provide noncooperative foundations for the hybrid cooperative/noncooperative game, using cheap-talk bargaining and axiomatic equilibrium selection. In Appendix Section B3 we explain how to generalize Miller and Watson's results to our setting with external enforcement.

Section I presents our general model and the definition of contractual equilibrium, expressed as a recursive characterization of equilibrium continuation values. This characterization extends the notion of *self-generation* (Abreu, Pearce, and Stacchetti 1990) for our contracting environment and is convenient for applications.

Section II presents our leading application, a principal-agent relationship with a costly and externally enforceable monitoring technology, which illustrates the components of our theory and all of our general results. In a setting with no verifiable information, optimal contracts are semi-stationary and specify mild monitoring for the current period but intense monitoring for future periods, which the players adjust each period in equilibrium. When we augment the example by adding a verifiable monitoring signal but no externally enforced contingent transfers, the optimal contract is no longer semi-stationary. But with contingent transfers, there is once again an optimal semi-stationary contract. This extension demonstrates the importance of contingent transfers for our main result: if there is verifiable information but externally enforced transfers are constrained by, say, limited liquidity or legal constraints, then semi-stationary contracts will not generally be optimal.

Section III presents our general results. Section IIIA shows how to construct optimal contracts within the class of semi-stationary contracts. Theorem 1, in Section IIIB, shows that semi-stationary contracts are optimal in contractual settings with *externally enforced contingent transfers*. Theorem 2, in Section IIIC, obtains the same result for contractual settings with *no verifiable information*. Section IIID explains why improvements in external enforcement are always complementary with self-enforcement in our model. Appendix A provides the proof of Theorem 1, and Appendix B provides foundations for contractual equilibrium and a discussion of technical issues related to existence.

In Section IV, we expand the application from Section IIB by allowing option contracts, in which one player verifiably selects from a menu of monitoring/payment pairs. This application shows how giving parties the ability to contractually allocate decision rights can expand the scope for cooperation. In this case, whether decision rights are optimally allocated to the manager or to the worker depends on their relative bargaining strengths.

I. The Model

We generalize the model of Miller and Watson (2013) by adding external enforcement. Players 1 and 2 play a relational contracting game in discrete time over an infinite horizon, with discount factor $\delta \in (0,1)$. In each period, there are two phases: the *negotiation phase*, followed by the *action phase*. In the negotiation phase, the players jointly decide to form or revise their contract and make an immediate monetary transfer. In the action phase, the players individually select actions in a stage game and receive payoffs. External enforcement is incorporated into the stage game, which may vary from period to period as specified by the players' contract. At the end of each period the players jointly observe an unverifiable draw from a randomization device that we assume is uniformly distributed on the unit interval. We normalize payoffs by multiplying by $1 - \delta$.

A. *Stage Games and External Contracts*

A *stage game* $\gamma = (A, X, \lambda, u, P)$, to be played in the action phase, has the following components:

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

Each player i takes an action $a_i \in A_i$. The action profile $a \in A$ determines the probability distribution $\lambda(a) \in \Delta X$ over outcomes. The realized outcome $x \in X$ is commonly observed by the players, but only the partition element that contains x, denoted $P(x)$, is verifiable. Though stage-game payoffs can in general depend on both the action profile a and the outcome x, we define $u(a)$ as the expected payoff over $x \sim \lambda(a)$ when the players choose action profile a. Player i observes only the outcome x and her own action a_i.[7]

In each period, the players' current external contract specifies a stage game for them to play in the action phase, as a function of the history of stage-game outcomes. Formally, there is a set Γ of feasible stage games, and we let $\mathcal{X} \equiv \cup\{X \mid (A, X, \lambda, u, P) \in \Gamma\}$ be the set of possible stage-game outcomes. Let $H^X \equiv \cup_{k=0}^{\infty} \mathcal{X}^k$ be the space of finite-length outcome histories, where $\mathcal{X}^0 \equiv \{h^0\}$ is the singleton consisting of the null history at the start of the game. An *external contract* is a function $c : H^X \to \Gamma$, where $c(h)$ is the stage game to be played in the period following outcome history $h \in H^X$. As noted in the introduction, we use the qualifier "external" to distinguish this from the self-enforced part of the players' contract, their coordinated play in the action phase over time. But where it would not cause confusion, we drop the qualifier and say simply "contract."

In our analysis, we study such contracts in the form of "continuation contracts." Given a history of outcomes through period $t - 1$, the continuation contract from period t gives the stage game in each period $\tau \geq t$ as a function of the history of outcomes from t until $\tau - 1$. The continuation contract may be interpreted as specifying (i) the stage game to be played in period t and (ii) a mapping from the stage-game outcome to the continuation contract in period $t + 1$. Formally, for any $c : H^X \to \Gamma$, let $g(c) \equiv c(h^0)$ be the stage game prescribed for the initial period. For any $x \in \mathcal{X}$ and $h \in H^X$, where h is k periods in length, let xh denote the $(k + 1)$-period outcome history in which x is followed by the sequence h. Define $c|x : H^X \to \Gamma$ by $(c|x)(h) \equiv c(xh)$ for every $h \in H^X$. If the players operate under continuation contract c^t in period t, then they play stage game $g(c^t)$ and, after realizing outcome x^t in period t, they will inherit continuation contract $c^t|x^t$ in period $t + 1$.

[7] To model a setting in which players observe each other's actions, X and λ can be defined so that the outcome reveals the action profile. The framework also allows for applications in which the players may not observe their own payoffs.

External contracts can depend only on information that is verifiable. This means the transition from a continuation contract in one period to the continuation contract in the following period must be measurable with respect to the partition of stage-game outcomes. Formally, for any contract c, and letting $(A, X, \lambda, u, P) = g(c)$, the contract *respects verifiability* if $x \in P(x')$ implies $c|x = c|x'$ for all $x, x' \in X$. Let C be the set of contracts that respect verifiability.[8]

B. *The Relational Contracting Game*

We now describe the relational contracting game. In each period t, players enter the negotiation phase with a contract $\hat{c}^t \in C$, inherited from period $t-1$. The inherited contract at the beginning of the game, denoted $c^0 \equiv \hat{c}^1$, is exogenous and represents the default legal rule. In the negotiation phase, the players bargain to select a contract $c^t \in C$ and an immediate monetary transfer $m^t \in \mathbb{R}^2_0$, where $\mathbb{R}^2_0 \equiv \{m \in \mathbb{R}^2 | m_1 + m_2 = 0\}$ is the set of balanced transfers. The negotiated transfer is enforced automatically with the agreement. If the players do not reach an agreement, then they operate under the inherited contract, so $c^t = \hat{c}^t$ and the transfer is zero.

We model interaction in the negotiation phase cooperatively. The bargaining protocol is represented by a fixed vector of bargaining weights $\pi = (\pi_1, \pi_2)$ satisfying $\pi_1, \pi_2 \geq 0$ and $\pi_1 + \pi_2 = 1$. The bargaining weights can be viewed as a reduced form of a noncooperative bargaining protocol, such as one in which π_i is the probability that player i gets to make an ultimatum offer. Appendix Section B3 discusses the connections between the cooperative and noncooperative approaches, along the lines of Miller and Watson (2013) and Watson (2013).

In the action phase of period t, the players simultaneously choose actions in the stage game $\gamma^t = (A^t, X^t, \lambda^t, u^t, P^t)$ prescribed by the current contract c^t. Action profile $a^t \in A^t$ leads to an outcome x^t, distributed according to $\lambda^t(a^t)$. Along with the outcome, the players observe the draw of the public randomization device.

The payoffs within period t are given by the sum of any monetary transfer and the stage-game payoffs, normalized by $1 - \delta$, so the expected payoff vector is $(1-\delta)(m^t + u^t(a^t))$. As the game progresses, the players' behavior and the outcomes of the exogenous random variables induce a sequence of transfers and stage-game payoffs, so the continuation payoff vector from any period τ is the expected value of

$$\sum_{t=\tau}^{\infty} \delta^{t-\tau}(1-\delta)\big(m^t + u^t(a^t)\big),$$

conditioned on the history prior to time τ and the specification of behavior from period τ.

In summary, the *contractual setting* is described by the set of feasible stage games Γ (and its associated set of contracts C that respect verifiability), the default contract c^0,

[8] Limitations on external enforcement can be modeled as restricting the players to a subset $\hat{C} \subset C$ of enforceable contracts. Our analysis applies without alteration if $\hat{C}$ is closed under the transition relation. In online Appendix Section C.3 we provide an existence result for finite $\hat{C}$. Otherwise we shall not constrain C.

and bargaining weights π. We make two regularity assumptions throughout. First, we assume that c^0 specifies the same stage game for every history and that this stage game has a Nash equilibrium. Second, we assume that Γ has uniformly bounded joint values: there is a number $\vartheta \in \mathbb{R}$ such that for every stage game $(A, X, \lambda, u, P) \in \Gamma$ and every $a \in A$, we have $-\vartheta \leq u_1(a) + u_2(a) \leq \vartheta$.

C. *Contractual Equilibrium Values*

We analyze behavior using the concept of *contractual equilibrium* (Miller and Watson 2013, Watson 2013), which requires the following: in the action phase, each player's individual action is optimal in response to the other player's action and the equilibrium specification of future behavior. In the negotiation phase, the players reach an agreement consistent with the generalized Nash bargaining solution with bargaining weights π, where the disagreement point entails equilibrium play from the action phase of the current period under the inherited contract with no immediate transfer. The players renegotiate their entire contract in the negotiation phase, including the external contract c, their coordinated play in the stage game of the current period, and their plans for how future play under disagreement depends on the history of stage-game outcomes. Thus, an agreement in one period implicitly specifies the disagreement points in future periods.[9]

There are two standard approaches to characterizing equilibria in repeated games. The first involves describing strategies for the dynamic game and then stating and evaluating equilibrium conditions on the strategy space. The second characterizes the set of equilibrium continuation values recursively, following Abreu, Pearce and Stacchetti (1990), with equilibrium conditions expressed through dynamic programming. While both approaches extend to contractual equilibrium, we follow the recursive approach for convenience. Appendix Section B1 exposits the strategic approach and the links between the two approaches.

Because long-term contracts render the relational contracting game nonstationary, the set of continuation values attainable from a given period depends on the inherited contract. We therefore deal with collections of the form $\mathcal{W} = \{W(c)\}_{c \in C}$ where, for every $c \in C$, $W(c) \subset \mathbb{R}^2$ is the set of equilibrium continuation values from the beginning of a period in which c is the inherited contract.[10] Our characterization of equilibrium values extends the notion of self-generation (Abreu, Pearce, and Stacchetti 1990), as we describe next.

Note that in a given period under contract c, the players interact in stage game $g(c) \equiv (A, X, \lambda, u, P)$ and will get an outcome $x \in X$, leading to inherited contract $c|x$ in the next period. The players will then anticipate coordinating on some continuation value in $W(c|x)$ in the next period. Since the players can randomize over continuation values by conditioning on the draw of the public-randomization device, they are essentially picking a value in the convex hull of $W(c|x)$, which we denote $\operatorname{co} W(c|x)$. Let $y(x)$ denote the expected continuation value that the players coordinate on in the event that outcome x occurs in the current period. Also, given

[9] As in perfect public equilibrium, contractual equilibrium assumes that the players' equilibrium behavior is conditioned only on their common history, so the bargaining set and disagreement point are commonly known.

[10] We need to allow $W(c) = \varnothing$ for technical reasons discussed in footnote 12 and Appendix Section B2.

such a continuation function $y : X \rightarrow \mathbb{R}^2$, let $\bar{y} : A \rightarrow \mathbb{R}^2$ be the expected continuation function $\bar{y}(a) \equiv E_x[y(x) | x \sim \lambda(a)]$.

Incorporating the anticipated continuation value, in the current period the players' interaction is effectively to play the induced static game

$$\langle A, (1-\delta)u(\cdot) + \delta\bar{y}(\cdot)\rangle, \tag{1}$$

where A is the set of action profiles, and payoffs are the convex combination of stage-game payoffs and continuation values. The players can self-enforce any mixed action profile $\alpha \in \Delta A$ that is a Nash equilibrium of this induced game, resulting in continuation value

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

from the action phase in the current period.[11]

DEFINITION 1: *Given $\gamma = (A, X, \lambda, u, P) \in \Gamma$ and $y : X \rightarrow \mathbb{R}^2$, call action profile $\alpha \in \Delta A$* ***enforced relative to*** *γ* ***and*** *y if it is a Nash equilibrium of Induced Game* 1.

DEFINITION 2: *Given $\mathcal{W} = \{W(c)\}_{c \in C}$, consider any contract $c \in C$, and let $g(c) = (A, X, \lambda, u, P)$. Say that $w \in \mathbb{R}^2$ is* ***c-supported relative to*** *$\mathcal{W}$ if there exist $\alpha \in \Delta A$ and $y : X \rightarrow \mathbb{R}^2$ such that $y(x) \in \operatorname{co} W(c|x)$ for all $x \in X$, α is enforced relative to $g(c)$ and y, and equation* (2) *holds.*

Turning to the negotiation phase of the current period, under inherited contract $\hat{c}$ the players would coordinate on some $\hat{c}$-supported continuation value $\underline{w}$ in the event that they fail to make an agreement. Thus, $\underline{w}$ is the disagreement point for negotiation in the current period. The Nash bargaining solution predicts that the players renegotiate to a contract c and coordinate on a c-supported continuation value that maximizes their joint value,

$$L(\mathcal{W}) \equiv \max\{w_1 + w_2 \,|\, c \in C \text{ and } w \text{ is } c\text{-supported relative to } \mathcal{W}\}, \tag{3}$$

and they make an immediate transfer to split the surplus in proportion to their bargaining weights. We call $L(\mathcal{W})$ the *level* of the collection. Because an equilibrium collection $\mathcal{W}$ gives the continuation values available from every period, it must satisfy the following self-generation condition.

DEFINITION 3: *Say that a collection $\mathcal{W} = \{W(c)\}_{c \in C}$ is* ***bargaining self-generating*** *(**BSG**) if for every $\hat{c} \in C$ and $w \in W(\hat{c})$, there exists a value $\underline{w}$ that is $\hat{c}$-supported relative to $\mathcal{W}$ such that $w = \underline{w} + \pi(L(\mathcal{W}) - \underline{w}_1 - \underline{w}_2)$.*

[11] Here ΔA is defined as the space of uncorrelated probability distributions over A.

The BSG condition captures the idea of *internal consistency* in that the bargaining solution selects among all continuation values attainable relative to $\mathcal{W}$. Contractual equilibrium incorporates the additional condition of *external consistency*, meaning that the players attain the maximum joint value over all internally consistent equilibria.

DEFINITION 4: *A collection $\mathcal{W}$ is called a* ***contractual equilibrium value*** (***CEV***) ***collection*** *if it is BSG and its level $L(\mathcal{W})$ is maximal among the set of BSG collections.*

We will say that *contractual equilibrium exists* if there is a CEV collection $\mathcal{W}$ with the property that $W(c^0) \neq \varnothing$. Existence of contractual equilibrium is analyzed in the context of our main characterization results in the next section.[12] At this point, we have the following immediate implication of the CEV definition.

LEMMA 1: *For a given contractual setting, all CEV collections attain the same level.*

For every $c \in C$, let $W^*(c)$ be the union of all $W(c)$ sets, over all CEV collections, and let $\mathcal{W}^* \equiv \{W^*(c)\}_{c \in C}$. Under conditions for existence developed in the next section, $\mathcal{W}^*$ is also a CEV collection and so we refer to it as the *maximal CEV collection*. We call c^* an *optimal contract* if it solves the maximization problem that defines $L(\mathcal{W}^*)$ in equation (3). We sometimes refer to the equilibrium level as L^*.

Clearly, from Lemma 1 and the BSG definition, we have $w_1 + w_2 = L^*$ for every c and every $w \in W^*(c)$. Also, for an arbitrary set $Y \subset \mathbb{R}^2$ of constant joint value, let us refer to the vertical/horizontal distance between its extreme points as its *span*:

$$\text{span}(Y) \equiv \sup\{w_1 - w_1' \,|\, w, w' \in Y\}.$$

We shall say that *Y attains its span* if it contains its extreme points, so there are elements $z^1, z^2 \in Y$ such that $z_1^2 - z_1^1 = \text{span}(Y)$.

As for which payoff vector in a CEV collection the players obtain from the start of the game, it depends on what their continuation play would be if they fail to agree in the first period. For instance, if the initial contract c^0 specifies a constant stage game that represents the players' outside values, and we normalize these outside values to zero, then in a contractual equilibrium the players get payoffs of exactly πL^*. That is, they split the surplus (using voluntary transfers in the negotiation phase) relative to their outside values in accordance with their bargaining weights.

II. Example: Choice of Monitoring Technology

This section analyzes a simple example, in several variations, to illustrate our approach and main results, in a way that we hope also provides some novel economic

[12] Existence of a BSG collection requires existence of a maximum in expression (3). Note that contractual equilibrium can exist with $W(c) = \varnothing$ for some values of c, which we allow to deal with off-path contracts under which there would be no best response in the action phase (see Appendix Section B2). Also, for convenience we allow $W(c)$ to be empty if c is a contract that would never be inherited. In online Appendix Section C.3 we prove an existence result for settings with a finite number of external contracts, where $W(c) \neq \varnothing$ for all c.

insights. Consider a relationship between a worker (player 1) and a manager (player 2), with an externally enforced monitoring technology. In the action phase, the players interact in a stage game parameterized by a monitoring level $\mu \in [0,1]$. The worker privately chooses effort $a_1 \in A_1 = \{0,1\}$. High effort, $a_1 = 1$, imposes a cost $\beta \in (0,1)$ on the worker and yields a benefit of 1 to the manager, both in monetary terms. The manager has no action but pays $k(\mu)$ for the monitoring technology. The stage-game payoff vector is therefore given by $u(a_1) = (-\beta a_1, a_1 - k(\mu))$. Assume $k(\cdot)$ is differentiable, $k' > 0$, and $\beta + k(1) \leq 1$.

The stage-game outcome $x \in X = \{1,0\}$ is a signal of the worker's effort choice. We call $x = 1$ the "high" signal and $x = 0$ the "low" signal. If the worker exerts high effort then the signal is high for sure, but if the worker exerts low effort then the signal is high with probability $1 - \mu$ and low with probability μ. The manager does not observe the worker's effort choice or the payoff he receives.[13] Assume that the signal is not verifiable (the external enforcer cannot distinguish between $x = 1$ and $x = 0$), and so $P = \{\{0,1\}\}$.

Because nothing is verifiable, an external contract is simply a sequence $c = \{\mu^\tau\}_{\tau=1}^{\infty}$, where μ^1 is the monitoring level specified for the current period, μ^2 is the monitoring level specified for the next period, and so on. Note that regardless of the outcome x in the current period, the contract inherited in the following period is $c|x = \{\mu^\tau\}_{\tau=2}^{\infty}$.

A. *Fixed Monitoring Technology*

As a benchmark, we first examine the setting in which the monitoring technology μ is exogenously fixed and constant over time. That is, Γ contains just one stage game, so in the negotiation phase, the players have only their immediate transfer and their self-enforced continuation play to discuss. There is just one set of continuation values to calculate, W, which we write without reference to the lone contract c^0.

This relationship falls within the class analyzed by Miller and Watson (2013), where the contractual equilibrium value set W^* is easily characterized.[14] Because every element of W^* has the same joint value L^*, W^* is a subset of a line segment of slope -1. In fact, W^* attains its span, and we let z^1 and z^2 denote the extreme points, where z^1 gives the worst continuation value for player 1 and z^2 gives the worst for player 2. Other points in W^* are inessential to the equilibrium construction because the players can utilize the public randomization device to coordinate on any point in the convex hull as an expected continuation value. Depending on parameter values, either high effort will be sustainable and $L^* = 1 - \beta - k(\mu)$, or high effort cannot be achieved and $L^* = -k(\mu)$.

[13] Alternatively, we could assume that the manager's payoff depends only on the monitoring signal, equaling 1 if $x = 1$ and $-(1 - \mu)/\mu$ if $x = 0$, which implies the same payoff function u.

[14] It is also easy to calculate, as a benchmark, the optimal perfect public equilibrium in a setting with no negotiation but still with voluntary transfers, as analyzed by Levin (2003). 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 $(1 - \delta)\beta \leq \delta\mu(1 - \beta)$. 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.

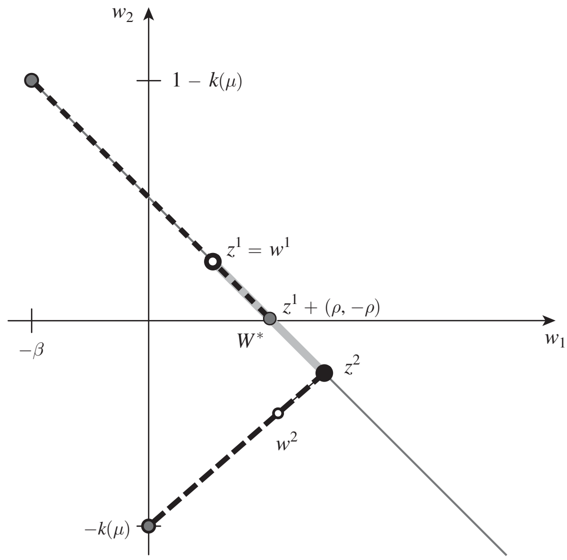

Figure 1. Contractual Equilibrium with Fixed Monitoring

Notes: Figures in Section II are drawn to scale using parameters $\beta = 1/4$, $k(\mu) = (3/4)\mu$, $\delta = 3/4$, and $\pi_1 = \pi_2 = 1/2$. Disagreement point $\underline{w}^1$ or any other point in W^* is attained by making an appropriate transfer, playing $a_1 = 1$, and continuing with promised utility $z^1 + (\rho, -\rho)$ if the signal is $x = 1$, but with z^1 if $x = 0$. Since $\underline{w}^1 \in W^*$, $z^1 = \underline{w}^1$. Disagreement point $\underline{w}^2$ is attained by playing $a_1 = 0$ and continuing with z^2 regardless of x. When $\underline{w}^2$ is the disagreement point, the parties negotiate to z^2.

Let us proceed under the presumption that $L^* = 1 - \beta - k(\mu)$. With reference to the BSG condition, we can determine z^1 and z^2 by characterizing the associated disagreement points $\underline{w}^1$ and $\underline{w}^2$ for which $z^1 = \underline{w}^1 + \pi(L^* - \underline{w}^1_1 - \underline{w}^1_2)$ and $z^2 = \underline{w}^2 + \pi(L^* - \underline{w}^2_1 - \underline{w}^2_2)$. Here $\underline{w}^1$ must be the supported continuation value from the action phase that is most favorable to player 2, whereas $\underline{w}^2$ is the one most favorable to player 1.

Disagreement point $\underline{w}^1$ is characterized as follows and displayed in Figure 1. The players coordinate on $a_1 = 1$ being played in the current period. Then if the signal is high, they coordinate to achieve expected continuation value $z^1 + (\rho, -\rho)$ from the next period. If the signal is low, they coordinate on z^1 from the next period. Thus,

$$\underline{w}^1 = (1 - \delta)\big(-\beta, 1 - k(\mu)\big) + \delta z^1 + \delta(\rho, -\rho). \tag{4}$$

The value of ρ must be large enough to ensure that the worker does not want to deviate to low effort, knowing that such a deviation would be detected with probability μ, and then punished:

$$-(1 - \delta)\beta + \delta\big(z^1_1 + \rho\big) \geq (1 - \delta) \cdot 0 + \mu\delta z^1_1 + (1 - \mu)\delta\big(z^1_1 + \rho\big).$$

This incentive constraint simplifies to $\mu\delta\rho \geq \beta(1 - \delta)$. Because we are characterizing the supported continuation value that is worst for player 1, it is optimal to pick the smallest possible value of ρ, so we set $\rho = (1 - \delta)\beta/\delta\mu$. Because play in the current period is efficient, $\underline{w}_1 + \underline{w}_2 = L^*$, so there is no bargaining surplus; thus $z^1 = \underline{w}^1$. Combining this with equation (4), inserting the values of ρ and L^*, and solving for z^1 yields

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

The payoffs reflect the worker's rent from exerting effort under imperfect monitoring.

Disagreement point $\underline{w}^2$ is characterized as follows and displayed in Figure 1. The players coordinate on $a_1 = 0$ being played in the current period and, regardless of the signal realization, they coordinate to achieve continuation value z^2 from the next period. Thus,

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

Combining this with $z^2 = \underline{w}^2 + \pi(L^* - \underline{w}_1^2 - \underline{w}_2^2)$ and inserting the value of L^*, we obtain

$$z^2 = (0, -k(\mu)) + \pi(1 - \beta).$$

Here the payoffs reflect that the parties share the bargaining surplus in proportion to their bargaining weights.

The final equilibrium condition is that $\rho \leq \text{span}(W^*)$; that is, the bonus in continuation value that the worker receives for a high signal must be attainable. Noting that

$$\text{span}(W^*) \equiv z_1^2 - z_1^1 = z_2^1 - z_2^2 = \pi_1(1 - \beta) - \left(\frac{\beta}{\mu} - \beta\right)$$

and recalling that $\rho = (1 - \delta)\beta/\delta\mu$, we find that the condition for sustaining high effort in contractual equilibrium simplifies to

$$\beta \leq \mu\delta(\pi_1 + \beta - \pi_1\beta). \tag{5}$$

If this inequality does not hold, then high effort cannot be sustained, the level is $L^* = -k(\mu)$, and $(0, -k(\mu))$ is the unique contractual-equilibrium value.[15]

It is important to note how the equilibrium 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. There is thus a trade-off in setting the monitoring level: a high enough span is needed for the worker's incentive condition, but it comes at a higher monitoring cost. The monitoring level that maximizes welfare is the lowest that satisfies condition (5), which is $\mu = \beta/(\delta\beta + \delta\pi_1 - \delta\pi_1\beta)$.

[15] Unless $\mu = \pi_1 = 1$, the condition for sustaining high effort in the contractual equilibrium is stricter than the corresponding condition for the optimal perfect public equilibrium described in footnote 14. The difference arises because the perfect public equilibrium employs punishments that would not survive renegotiation.

B. *Contractible Monitoring Technology*

Now suppose that Γ contains all monitoring levels $\mu \in [0,1]$, so the players can write a contract that specifies any sequence $c = \{\mu^\tau\}_{\tau=1}^{\infty}$. For any contract $c \in C$, let $z^1(c)$ and $z^2(c)$ denote the extreme points of $W^*(c)$, which attains its span as in the previous setting. As before, we let $z^i(c)$ denote the worst point for player i.

It turns out that, in contractual equilibrium, stationary contracts (specifying the same μ in all periods) are suboptimal in the present setting. Instead, the optimal contract is *semi-stationary*, specifying one monitoring level $\hat{\mu}$ for the current period and another level $\underline{\mu}$ for all future periods. Then in equilibrium the inherited contract is always $\{\underline{\mu}\}_{t=1}^{\infty}$, and the players always renegotiate to specify $\hat{\mu}$ for the current period and $\underline{\mu}$ 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, which supports wide-ranging disagreement points. The players anticipate renegotiating in the future to lower the monitoring level one period at a time. Renegotiation shifts every disagreement point to a continuation value in the direction of π because players share surplus in this proportion, so renegotiation ensures a high joint value while maintaining the large span of continuation values.

To perform the analysis formally and to calculate the monitoring levels $\hat{\mu}$ and $\underline{\mu}$ that are featured in the optimal contract, take as given any contract $c = \{\mu^\tau\}_{\tau=1}^{\infty}$ and let $c' = c|x = \{\mu^\tau\}_{\tau=2}^{\infty}$ denote the inherited contract in the next period. We shall express $z^1(c)$ and $z^2(c)$ as functions of $z^1(c')$ and $z^2(c')$ which, in particular, relates $\mathrm{span}(W^*(c))$ to $\mathrm{span}(W^*(c'))$ and also helps us calculate L^*.

The specifications of disagreement play that support extreme points $z^1(c)$ and $z^2(c)$ are exactly as in the fixed-μ case, except that the continuation values in the following period are taken from the set $W^*(c')$. In the disagreement point associated with $z^1(c)$, players coordinate on play of $a_1 = 1$ in the current period and on continuation value $z^1(c') + x(\rho, -\rho)$ from the next period (giving a bonus of ρ to the worker if the signal is high):

$$(6) \qquad \underline{w}^1(c) = (1-\delta)\big(-\beta, 1 - k(\mu^1)\big) + \delta z^1(c') + \delta(\rho, -\rho).$$

Since the last term is a transfer and $z^1(c')$ has joint value L^*, the negotiation surplus derives entirely from changing the monitoring level in the current period, and we have $z^1(c) = \underline{w}^1(c) + (1-\delta)\pi\big(L^* - (1 - \beta - k(\mu^1))\big)$. Combining this with equation (6) yields

$$(7) \qquad z^1(c) = (1-\delta)\left(\frac{\beta}{\mu^1} - \beta, 1 - \frac{\beta}{\mu^1} - k(\mu^1)\right) + (1-\delta)\pi\big(L^* - 1 + \beta + k(\mu^1)\big) + \delta z^1(c'),$$

where we have set $\rho = (1-\delta)\beta/\delta\mu^1$ to make the worker's incentive constraint bind.[16] The payoff to the worker in the current period reflects her rent from effort plus her share of the negotiation surplus.

The disagreement point associated with $z^2(c)$, as before, entails play of $a_1 = 0$ and coordination on continuation value $z^2(c')$ from the next period, implying

$$\underline{w}^2(c) = (1-\delta)\big(0, -k(\mu^1)\big) + \delta z^2(c').$$

The bargaining solution implies $z^2(c) = \underline{w}^2(c) + (1-\delta)\pi(L^* + k(\mu^1))$. Combining these expressions yields

$$z^2(c) = (1-\delta)\big(0, -k(\mu^1)\big) + (1-\delta)\pi\big(L^* + k(\mu^1)\big) + \delta z^2(c'). \tag{8}$$

Recalling the definition of span, we subtract equation (7) from equation (8) to obtain

$$\text{span}\big(W^*(c)\big) = (1-\delta)(1-\beta)\pi_1 - (1-\delta)\beta\,\frac{1-\mu^1}{\mu^1} + \delta\,\text{span}\big(W^*(c')\big). \tag{9}$$

Suppose that we want to design a contract to maximize $\text{span}(W^*(c))$. Because expression (9) is increasing in μ^1 and in $\text{span}(W^*(c'))$, we should set $\mu^1 = 1$ and, by induction, specify the same maximal monitoring level in all future periods. Therefore, the span is maximized by the contract $\underline{c} \equiv \{1\}_{\tau=1}^{\infty}$. Inserting $c = c' = \underline{c}$ into expression (9) and simplifying yields $\text{span}(W^*(\underline{c})) = \pi_1(1-\beta)$, which is strictly higher than the span of W^* in the fixed-μ setting. Correspondingly, the sufficient condition for enforcing high effort, $\rho \leq \text{span}(W^*(\underline{c}))$, is weaker than inequality (5).

Of course, when the players negotiate in a given period, they will want to maximize the span not from the current period but from the *next* period, which allows them to support high effort in the current period at the lowest possible monitoring level (to save on monitoring costs that they will actually have to pay). Therefore they should agree on a contract that makes $\underline{c}$ the inherited contract in the next period. To calculate the monitoring level needed to support high effort in the current period, recall that the worker must be rewarded for high output with a bonus in continuation value of at least $(1-\delta)\beta/\delta\mu$, where μ is the monitoring level in the current period. The best choice for μ is the smallest value that satisfies the constraint $(1-\delta)\beta/\delta\mu \leq \text{span}(W^*(\underline{c}))$, which is

$$\hat{\mu} = \frac{(1-\delta)\beta}{\pi_1\delta(1-\beta)}. \tag{10}$$

To summarize, in the contractual equilibrium the players initially choose contract $c^* = \{\mu^\tau\}_{\tau=1}^{\infty}$ defined by $\mu^1 = \hat{\mu}$ and $\mu^\tau = 1$ for $\tau = 2, 3, \ldots$. In each subsequent period, the players inherit contract $\underline{c}$ and renegotiate back to c^*. In other words, they revise their inherited contract by specifying $\hat{\mu}$ in the current period but

[16] If $(1-\delta)\beta/\delta\mu^1 > \text{span}\big(W^*(c')\big)$ then high effort cannot be supported in disagreement and $z^1(c)$ is the same as $z^2(c)$ characterized below.

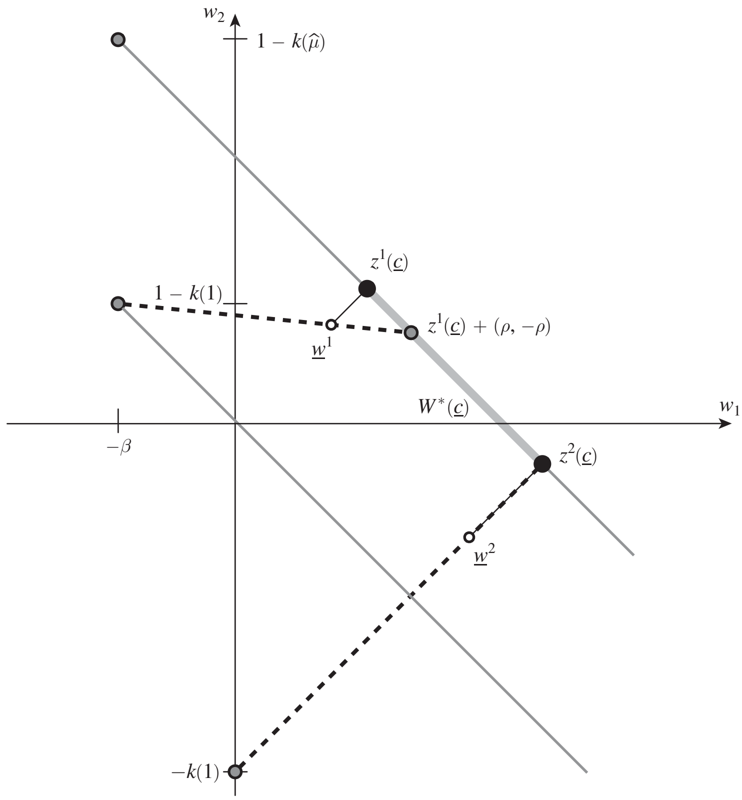

Figure 2. Contractual Equilibrium with Contractible Monitoring Technology

Notes: Disagreement point $\underline{w}^1$ is attained by playing $a_1 = 1$ under contract $\underline{c}$ (with monitoring level $\underline{\mu} = 1$) and continuing with promised utility $z^1(\underline{c}) + (\rho, -\rho)$ if the signal is $x = 1$, but with $z^1(\underline{c})$ if $x = 0$. Disagreement point $\underline{w}^2$ is attained by playing $a_1 = 0$ under $\underline{c}$ and continuing with $z^2(\underline{c})$ regardless of x. When $\underline{w}^i$ is the disagreement point, the parties negotiate to contract c^* (with monitoring level $\hat{\mu}$) and utility $z^i(\underline{c})$. Any point in W^* is attained by playing $a_1 = 1$ under c^*, making an appropriate transfer, and continuing with $z^2(\underline{c})$ if $x = 1$, but with $z^1(\underline{c})$ if $x = 0$.

leave the specified monitoring level at 1 for all future periods. The equilibrium continuation values and disagreement points are displayed in Figure 2. It is easy to verify that $\hat{\mu}$ is strictly less than the optimal monitoring level in the fixed-μ setting, for parameter values under which cooperation can be sustained. Therefore, the players get a strictly higher joint value from the optimal semi-stationary contract than from the best stationary contract.

C. *Verifiable Signal*

The example presented in the previous subsections illustrates one of our general results: semi-stationary contracts are optimal in settings with no verifiable information. We next show that this result does not extend to all settings with verifiable information. To do this, we examine an extension of the example in which the stage-game outcome x is verifiable. As before, external enforcement entails only operation of the monitoring technology at the contractually specified level, but now the sequence of monitoring levels can be conditioned on past realizations of x. We assume that externally enforced outcome-contingent transfers are not available; such transfers are discussed in the next subsection.

A semi-stationary contract does not condition the monitoring level on past realizations of x; it specifies a monitoring level $\hat{\mu}$ in the current period and a monitoring level $\underline{\mu}$ for all future periods *regardless* of the history of signal realizations. The best semi-stationary contract c^* is exactly as described in the previous subsection, with $\hat{\mu}$ given by equation (10) and $\underline{\mu} = 1$. Players coordinate on behavior and continuation values as before.

But c^* is no longer optimal. To see why, recall that the worker will select high effort only if the difference between his continuation values following high and low signals is at least $(1-\delta)\beta/\delta\mu$, where μ is the monitoring level in the current period. Maximizing the difference allows μ, and hence the cost of monitoring, to be minimized. In the initial example, regardless of the contract c, these continuation values were required to be elements of a single set $W^*(c|x)$, because $c|x$ could not depend on x (it was unverifiable). However, with x now verifiable, the inherited contracts $c|1$ and $c|0$ may differ. Rewards and punishments may be enhanced by conditioning $c|x$ on x.

Specifically, to reward the worker following a high monitoring signal $(x = 1)$ in the current period, the inherited contract $c|1$ should maximize $z_1^2(c')$ over $c' \in C$, and the players should coordinate on $z^2(c|1)$, the worker's best continuation value. As in the initial example, $z_1^2(\cdot)$ is maximized by contract $\underline{c}$ specifying $\mu = 1$ in all periods regardless of the signal realizations.[17] Likewise, to best punish the worker after a low monitoring signal $(x = 0)$ in the current period, $c|0$ should minimize $z_1^1(c'')$ over $c'' \in C$, and the players should coordinate on $z^1(c|0)$. But $\underline{c}$ generally does not minimize $z_1^1(\cdot)$. In fact, a stationary contract specifying a lower monitoring level may be better, depending on cost parameters. Consider such a contract $\tilde{c}$ that specifies $\mu = \tilde{\mu} < 1$ in all periods regardless of the signal realizations. The disagreement point associated with $z^1(\tilde{c})$ involves high effort, as in the initial example. We find that if the marginal monitoring cost is sufficiently large at the maximal level (specifically, if $\pi_1 k'(1) > \beta$), then $z_1^1(\tilde{c})$ is increasing in $\tilde{\mu}$ for $\tilde{\mu}$ near 1. This implies that there is a value $\tilde{\mu} < 1$ for which $z_1^1(\tilde{c}) < z_1^1(\underline{c})$.[18]

Suppose we design a contract c so that $c|1 = \underline{c}$ and $c|0 = \tilde{c}$, where the players would coordinate on continuation value $z^2(\underline{c})$ following the high signal and $z^1(\tilde{c})$ following the low signal. Such a contract supports high effort in the current period with less monitoring than c^* requires, because $z_1^2(\underline{c}) - z_1^1(\tilde{c}) > \text{span}(W^*(\underline{c}))$. The associated continuation values and disagreement points are displayed in online Appendix Section C.4. Importantly, c is not semi-stationary because the monitoring level specified for the following period depends on the verifiable outcome in the

[17] Assuming that the maximum is attained, let c' maximize $z_1^2(\cdot)$. Clearly the disagreement point that favors player 1 involves low effort in the current period and continuation value $z_1^2(c')$ regardless of the signal, so equation (8) is valid with $c = c'$. Player 1's payoff is increasing in μ^1, implying $c' = \underline{c}$.

[18] Equation (7) is valid with $c = c' = \tilde{c}$ and $\mu^1 = \tilde{\mu}$ if $\tilde{\mu}$ is close to 1, because the disagreement behavior associated with $z^1(\tilde{c})$ is as described in the initial example. Algebra yields $z^1(\tilde{c}) = (\beta/\tilde{\mu} - \beta, 1 - \beta/\tilde{\mu} - k(\tilde{\mu})) + \pi(L^* - 1 + \beta + k(\tilde{\mu}))$. Increasing $\tilde{\mu}$ reduces the worker's rent from effort (the first term) but increases the negotiation surplus (the second term). For $\tilde{\mu}$ near 1, the latter dominates if $\pi_1 k'(1) > \beta$.

current period. The equilibrium level strictly exceeds what can be achieved by any semi-stationary contract.

D. *Contingent Transfers*

In this subsection, we preview our main result, that semi-stationary contracts are optimal in settings with *contingent transfers* (external enforcement of arbitrary budget-balanced monetary transfers as a function of the verifiable outcome). Note first that adding contingent transfers to the example discussed in the previous subsection, where x is verifiable, enables the moral-hazard problem to be solved without any relational incentives. It suffices to choose a stationary contract that specifies a large monetary bonus for the worker in the event of $x = 1$. However, the prospect of being forced to pay a large bonus could tempt the manager to manipulate the signal. To better illustrate our main result, we extend the example to allow for nonverifiable signal manipulation by the manager. We show that the contractual equilibrium is semi-stationary, with incentives for effort provided by contingent bonuses, and incentives to abstain from manipulation provided by self-enforcement.

We augment the example so that the manager can take an unverifiable action that costlessly "jams the signal." The manager's action in the stage game is denoted $a_2 \in A_2 = \{0,1\}$, where $a_2 = 0$ refers to jamming the signal and $a_2 = 1$ means not jamming it. Stage-game payoffs are the same as before; they do not depend on a_2. The outcome is now written $x = (x_1, x_2) \in \{0,1\} \times \{0,1\}$, where x_1 is the signal realization and $x_2 = a_2$. If the manager chooses $a_2 = 1$ then x_1 depends on a_1 and μ exactly as in the initial example. If the manager chooses $a_2 = 0$ then with probability ε the signal is jammed and $x_1 = 0$ regardless of the worker's action, and with probability $1 - \varepsilon$ the signal realization depends on a_1 and μ as before. The probability ε is a fixed parameter. Note that x_1 is verifiable, as in the previous subsection, while x_2 is not verifiable.

Contingent transfers are incorporated as follows. The external contract can specify, in addition to the monitoring level, a monetary transfer from the manager to the worker as a function of the verifiable x_1. Let $b_1(x_1) \in \mathbb{R}$ denote the transfer in the event of signal realization x_1. The set of stage games Γ is parameterized by $(\mu, b_1(1), b_1(0))$ and stage-game payoffs include the expected transfer as a function of the action profile. A semi-stationary contract specifies two combinations of a monitoring level and contingent transfers: $(\hat{\mu}, \hat{b}_1(1), \hat{b}_1(0))$ in force for the current period, and $(\underline{\mu}, \underline{b}_1(1), \underline{b}_1(0))$ in force for all future periods irrespective of signal realizations.

Such a contract turns out to be optimal. A key idea is that, because renegotiation ensures that all continuation values in $\mathcal{W}^*$ have the same joint value, shifts between them are equivalent to monetary transfers. So rather than having external enforcement of current-period actions occur through the inherited contract in the next period, by specifying $c|(1,x_2) \neq c|(0,x_2)$ so that $W^*(c|(1,x_2)) \neq W^*(c|(0,x_2))$, it could alternatively occur with a monetary transfer in the current period. This is possible because the continuation contract $c|x$ and the transfer $b_1(x)$ are both conditioned on only the verifiable signal x_1.

A complication arises, however, because players also rely on self-enforcement (coordinating on continuation values within each set $W^*(c|x)$), and generally they can condition their play on elements of the outcome that are unverifiable, in particular the manager's action a_2. There is no way to substitute for this using externally enforced transfers. But transfers can substitute for shifting from one *set* of continuation values $W^*(c|(1,x_2))$ to another set $W^*(c|(0,x_2))$, as long as the latter set has as much scope for enforcing actions in the current period as does the former. Self-enforcement is best served by a large span of continuation values, so, with appropriate transfers, it is optimal to specify $c|(1,x_2) = c|(0,x_2)$ and to let this be the contract with the largest span.

The foregoing logic is a key element in the proof of our main result. While the general analysis requires additional technical steps, it is straightforward to summarize the equilibrium characterization in the example; a few of the calculations are shown in online Appendix Section C.5. Using the same steps as in the initial example, we find that the largest span is achieved by the contract $\underline{c}$ that specifies $\mu = 1$, $b_1(1) = \beta$, and $b_1(0) = 0$ in all periods regardless of the stage-game outcomes. We obtain $\text{span}(W^*(\underline{c})) = \pi_1(1-\beta)$ as before, but now the disagreement point associated with $z^1(\underline{c})$ requires that $a = (1,1)$ be enforced. The optimal contract c^* satisfies $c^*|x = \underline{c}$ for all x, and it provides incentives to the worker through a current-period monetary bonus $b_1(1) - b_1(0) > 0$. The players coordinate on the manager's favorite continuation value $z^1(\underline{c})$ if $x_2 = 1$ (no jamming), and on $z^2(\underline{c})$ if $x_2 = 0$. Adjusting for the expected transfer, this provides incentives to the manager.

The manager's ability to jam the signal constrains the use of contingent transfers, but nonetheless L^* is higher than in the initial example. In fact, among the examples in this section, those with greater scope for external enforcement exhibit higher equilibrium welfare levels, illustrating the general complementarity result we derive in the next section.[19] To our main point, the optimal contract in this example is semi-stationary, specifying $\mu < 1$ and a transfer bonus $b_1(1) - b_1(0) = \beta/\mu$ in the current period, and specifying $\mu = 1$ and $b_1(1) - b_1(0) = \beta$ in all future periods regardless of the stage-game outcomes.

III. Optimal Contracts and Semi-Stationarity

This section develops our main results, which show that the findings in our leading example regarding semi-stationary contracts hold broadly. We begin with these general definitions.

DEFINITION 5: *A contract $c \in C$ is* ***stationary*** *if $c|x = c$ for every $x \in \mathcal{X}$.*

DEFINITION 6: *A contract $c \in C$ is* ***semi-stationary*** *if there is a stationary contract $\underline{c}$ such that $c|x = \underline{c}$ for all $x \in \mathcal{X}$. In this case, we say that c* ***transitions to*** *$\underline{c}$.*

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

[19] Adding signal jamming to Sections IIB and IIC would not affect equilibrium welfare levels, because those examples do not have externally enforced contingent transfers.

contract c starts with stage game $g(c)$ and then specifies $g(\underline{c})$ in all future periods regardless of the history.

The first subsection below provides an algorithm to find an optimal contract in an artificial setting in which the players are restricted to semi-stationary contracts. In the subsections that follow, Theorem 1 establishes that semi-stationarity is indeed optimal in contractual settings with externally enforced contingent transfers, provided that the algorithm has a solution, and Theorem 2 obtains the same result for contractual settings with no verifiable information. The algorithm can then be used to calculate an optimal contract. The last subsection explains why external enforcement and self-enforcement are always complementary.

A. *Optimization within the Class of Semi-Stationary Contracts*

We introduce two optimization problems that jointly identify a contract that attains the maximal level among semi-stationary contracts. The first optimization problem determines the stationary part of the contract by finding the maximal span of continuation values that can be supported in the current period, as a function of the span of continuation values in the next period. This exercise corresponds to the analysis behind equation (9) in the example in Section IIB. The second optimization problem maximizes the joint payoff attained in the current period, assuming that the span of continuation values in the next period is the maximal fixed point from the first problem. It corresponds to the analysis behind equation (10) in the example.

Because negotiation always leads to the same welfare level, in both optimization problems we normalize the continuation values from the action phase so that they lie on the line $\mathbb{R}^2_0$ with zero joint value. The normalization is done by shifting stage-game payoffs along a ray in the direction of relative bargaining powers, π, which translates a payoff vector u to the point $u - \pi(u_1 + u_2)$. (Intuitively, this corresponds to a bargaining outcome $u + \pi(L - u_1 - u_2)$ with L normalized to 0.) Likewise, we normalize expected continuation values from the next period to be on the line segment $\mathbb{R}^2_0(d) \equiv \{m \in \mathbb{R}^2 | m_1 + m_2 = 0 \text{ and } m_1 \in [0,d]\}$, for a given span d.

In the first optimization problem, to maximize the span for the current period we look for a stage game $\gamma = (A, X, \lambda, u, P)$ 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 action profile $\alpha \in \Delta A$ and continuation value function $y : X \to \mathbb{R}^2_0(d)$, define

$$\omega(\gamma, \alpha, y) = (1 - \delta)\big(u(\alpha) - \pi\big(u_1(\alpha) + u_2(\alpha)\big)\big) + \delta \bar{y}(\alpha).$$

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

$$\Lambda(d) \equiv \max \omega_1\big(\gamma, \alpha^2, y^2\big) - \omega_1\big(\gamma, \alpha^1, y^1\big), \tag{11}$$

by choice of

$$\gamma = (A, X, \lambda, u, P) \in \Gamma, \quad y^1, y^2 : X \to \mathbb{R}_0^2(d), \quad \text{and} \quad \alpha^1, \alpha^2 \in \Delta A,$$

subject to

α^1 is enforced relative to γ and y^1, and α^2 is enforced relative to γ and y^2.

For the second optimization problem, let $\Xi(d)$ denote the maximized joint payoffs, by choice of the stage game, enforced action profile, and continuation value function:

$$\Xi(d) \equiv \max u_1(\alpha) + u_2(\alpha), \tag{12}$$

by choice of

$$\gamma = (A, X, \lambda, u, P) \in \Gamma, \quad y : X \to \mathbb{R}_0^2(d), \quad \text{and} \quad \alpha \in \Delta A,$$

subject to

α is enforced relative to γ and y.

Assume that $\Lambda(d)$ is defined for all d and has a largest fixed point, denoted d^*, and that $\Xi(d^*)$ exists. Let $\underline{\gamma} = (\underline{A}, \underline{X}, \underline{\lambda}, \underline{u}, \underline{P})$, y^1, y^2, α^1, and α^2 denote any solution to optimization problem (11) for Λ evaluated at d^*. Let $\gamma^* = (A^*, X^*, \lambda^*, u^*, P^*)$, y^*, and α^* denote any solution to optimization problem (12) for Ξ evaluated at d^*, so $\Xi(d^*)$ is the maximum value. Define $\underline{c}$ to be the stationary contract that specifies stage game $\underline{\gamma}$ in every period, and define c^* to be the semi-stationary contract that specifies stage game γ^* for the current period and then transitions to $\underline{c}$.

If players were restricted to semi-stationary contracts, c^* would be optimal and the equilibrium level L^* would equal $\Xi(d^*)$. Further, there would be a CEV collection in which $W(\underline{c}) = \{z^1(\underline{c}), z^2(\underline{c})\}$ where, for $j = 1, 2$, the disagreement point is

$$\underline{w}^j = (1 - \delta)\underline{u}(\alpha^j) + \delta\big(z^1(\underline{c}) + \bar{y}^j(\alpha^j)\big),$$

and the bargaining solution implies $z^j(\underline{c}) = \underline{w}^j + \pi(L^* - \underline{w}_1^j - \underline{w}_2^j)$. Using these expressions, the definition of ω, and that span of $W(\underline{c})$ is d^*, we derive

$$z^1(\underline{c}) = \omega(\underline{\gamma}, \alpha^1, y^1) + \pi(1 - \delta)L^* + \delta z^1(\underline{c});$$

$$z^2(\underline{c}) = \omega(\underline{\gamma}, \alpha^2, y^2) + \pi(1 - \delta)L^* + \delta z^2(\underline{c}) + \delta(-d^*, d^*).$$

These correspond to equations (7) and (8) in the initial example.[20] Collecting the $z^1(\underline{c})$ and $z^2(\underline{c})$ terms gives a direct expression of these values.

[20] Specifically, substituting for ω we see that $z^j(\underline{c})$ has a current-period component that reflects the players' sharing of the bargaining surplus, and a next-period component that consists of the worst continuation value for

Although an optimal semi-stationary contract is meant to be renegotiated every period, even if no deviation occurred previously, we do not claim that such "on-path renegotiation" should be seen in reality. In fact, in an enriched model that allows the players to send a joint, verifiable message in the negotiation phase, an optimal contract can include a provision that renews the equivalent of c^* if the players issue a joint statement of confirmation. For example, many real contracts specify that terms can be renewed by mutual agreement.[21] Then, rather than having to renegotiate the entire contract, the players can negotiate to exercise the joint renewal option and make an associated transfer.[22] We have left renewal options out of our model for simplicity, and to highlight the intertemporal changes in operative contract terms that occur in equilibrium.

B. *Semi-Stationarity with Contingent Transfers*

Many settings allow for external enforcement of arbitrary budget-balanced transfers as a function of verifiable outcomes. Our main result is that semi-stationary contracts are optimal in contractual settings with such contingent transfers (under some technical conditions sufficient for existence, namely that optimization problems (11) and (12) have solutions). The algorithm developed in the previous subsection can then be used to find an optimal contract.

To see how we can describe external enforcement of contingent transfers, suppose the players want to write a contract that augments stage game $(A, X, \lambda, u, P) \in \Gamma$ with a budget-balanced, P-measurable transfer function $b : X \to \mathbb{R}^2_0$ that requires player 2 to pay player 1 a transfer of $b_1(x)$ when outcome x occurs. Let $\bar{b}(a) \equiv E_x[b(x) | x \sim \lambda(a)]$ be the expected transfer given action profile $a \in A$. The availability of this contingent transfer is equivalent to assuming that the stage game $(A, X, \lambda, u + \bar{b}, P)$ is included in Γ, where $u + \bar{b} : A \to \mathbb{R}^2$ is the new payoff function that incorporates the transfers.

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

Our main result is the following.

THEOREM 1: *Suppose the contractual setting has externally enforced contingent transfers. If optimization problems* (11) *and* (12) *have solutions for all* $d \geq 0$ *then*

player 1, $z^1(\underline{c})$, plus a transfer from player 2 to player 1, $\bar{y}^j(\alpha^j)$.

[21] One common phrasing is "subject to unlimited successive renewals upon mutual consent of the parties" (see, for example, the *Law Insider* database of contracts from SEC filings, at lawinsider.com). In law and economics, contract renewal has mainly been viewed through the lens of the hold-up problem: see Blair and Lafontaine (2005) on franchising; Dalen, Moen, and Riis (2006) on procurement; and Narasimhan (1989) for a legal perspective.

[22] If the players do not send the joint statement of confirmation, and if they do not renegotiate the contract entirely, then the contract would implement the equivalent of $\underline{c}$. Either party can trigger $\underline{c}$ by blocking any joint action, such as in response to the other party's refusal to agree to a transfer. An alternative enrichment would involve adding a round of verifiable messages and verifiable voluntary transfers prior to negotiation in each period, whereby coordinated messages and transfers would be interpreted as exercising joint options. We conjecture that an optimal contract that avoids renegotiation would exist, but this is a topic for future study.

contractual equilibrium exists and there is a semi-stationary optimal contract c^*. *The level is* $L^* = \Xi(d^*)$, *where* d^* *is the largest fixed point of* Λ (*which exists*).

We provide a heuristic argument here, assuming the existence of a CEV collection with certain properties. This argument expands on the logic described at the end of Section IID. The formal proof, in Appendix Section A, follows a different logical path that also establishes the existence of a CEV collection.

Suppose there exists a contractual equilibrium and there is a contract $\tilde{c}$ whose value set $W^*(\tilde{c})$ has the greatest span in the maximal CEV collection $\mathcal{W}^*$. By definition, there is an optimal contract $c^{**} \in C$, but it may not be semi-stationary. Using externally enforced transfers, we will construct another optimal contract c^* that is semi-stationary.

First, we will construct a stationary contract $\underline{c}$ from $\tilde{c}$, with the property that $W^*(\underline{c}) = W^*(\tilde{c})$. Let $(A, X, \lambda, u, P) = g(\tilde{c})$ be the stage game specified by $\tilde{c}$. By definition of bargaining self-generation, any continuation value $w \in W^*(\tilde{c})$ is the Nash bargaining solution relative to some disagreement point $\underline{w}$ that is $\tilde{c}$-supported relative to $\mathcal{W}^*$. We construct a contract $\underline{c}$ with the property that any $\tilde{c}$-supported disagreement point $\underline{w}$ is also supported by $\underline{c}$, where $\underline{c}$ uses contingent transfers rather than variations in continuation-value sets.

Because all continuation values are at the same level L^*, variations in convex sets of continuation values act essentially as transfers. Therefore, if contract $\tilde{c}$ calls for the next-period value set $W^*(\tilde{c}|x)$ to differ from $W^*(\tilde{c})$ for some outcome x, we can construct $\underline{c}$ to instead specify an externally enforced, budget balanced transfer $b(x) = (1, -1)(\delta/(1 - \delta))(z_1^1(\tilde{c}|x) - z_1^1(\tilde{c}))$ in the current period and specify $\underline{c}|x = \tilde{c}$, without disrupting any incentives in the stage game. There are two key elements of this construction. First, because the continuation contract mapping $\tilde{c}|\cdot$ is P-measurable, so is the transfer function b. Second, $\text{span}(W^*(\tilde{c})) \geq \text{span}(W^*(\tilde{c}|x))$, so self-enforcement is no more constrained by contract $\underline{c}$ than by $\tilde{c}$.[23] Further, because this construction implies $W^*(\underline{c}) = W^*(\tilde{c})$, we can modify $\underline{c}$ to specify $\underline{c}|x = \underline{c}$. We have thus constructed a stationary contract $\underline{c}$ with the desired property.

Next we construct our semi-stationary contract c^* from c^{**}. Using the same steps as above, we now let $(A, X, \lambda, u, P) = g(c^{**})$ and we specify a transfer of $b(x) = (1, -1)(\delta/(1 - \delta))(z_1^1(c^{**}|x) - z_1^1(\tilde{c}))$. Define c^* to be the semi-stationary contract that specifies the stage game $(A, X, \lambda, u + \bar{b}, P)$ in the current period and then transitions to $\underline{c}$. Since it enforces the same actions as c^{**} does, c^* also supports a continuation value at level L^* and is thus optimal.

The theorem provides sufficient conditions for existence of a CEV collection in terms of whether the optimization problems defining $\Lambda(d)$ and $\Xi(d)$ have solutions for all $d \geq 0$. Sufficient conditions for existence that can be expressed more directly on the primitives have eluded us. Appendix Section B2 illustrates some of the difficulties.[24] We expect, however, that the optimization problems defining $\Lambda(d)$

[23] Compared to the continuation function y that was used to $\tilde{c}$-support $\underline{w}$, to $\underline{c}$-support $\underline{w}$ we use continuation function y' given by $y'(x) = y(x) + (z_1^1(\tilde{c}|x) - z_1^1(\tilde{c}))(1, -1)$, where $y'(x) \in \text{co}\, W^*(\tilde{c})$ follows from $y(x) \in W^*(\tilde{c}|x)$ and $\text{span}(W^*(\tilde{c})) \geq \text{span}(W^*(\tilde{c}|x))$.

[24] In online Appendix Section C.3 we prove that a CEV collection exists if C and Γ are finite and every stage game is finite, but these conditions rule out contingent transfers. One might speculate that a CEV collection should

and $\Xi(d)$ can be evaluated for relevant applications, as the next section illustrates. In any case, constructing an optimal contract with contingent transfers still involves computing d^* from $\Lambda(d)$ and then solving $\Xi(d^*)$.

C. *Semi-Stationarity with No Verifiable Information*

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

DEFINITION 8: *The contractual setting has* ***no verifiable information*** *if for every stage game* $\gamma = (A, X, \lambda, u, P) \in \Gamma$, *the partition P is trivial:* $P = \{X\}$.

Without verifiable information, a contract c can specify the sequence of stage games to be played but cannot condition the sequence on the history of stage-game outcomes. For instance, the initial example in Section II has no verifiable information, because the external enforcer cannot verify the monitoring signal. The following result shows that semi-stationarity is optimal in such settings.

THEOREM 2: *Suppose the contractual setting has no verifiable information. If optimization problems* (11) *and* (12) *have solutions for all* $d \geq 0$ *then contractual equilibrium exists and there is an optimal contract* c^* *that is semi-stationary. The level is* $L^* = \Xi(d^*)$, *where* d^* *is the largest fixed point of* Λ (*which exists*).

PROOF:

We prove this theorem by transforming the contracting environment into one to which Theorem 1 applies. For any relational contract setting, augment Γ so that there are externally enforced contingent transfers. This will change neither the CEV collections nor optimization problems (11) and (12), because the absence of verifiable information means that only a constant transfer can be specified in any period, and the players can already achieve such a transfer in the course of bargaining. From Theorem 1, we know contractual equilibrium exists and there is a semi-stationary optimal contract. If this contract specifies selection of nonzero externally enforced transfers, simply replace these with transfers in the bargaining phase and the equilibrium conditions remain satisfied. ■

D. *Complementarity of External Enforcement and Self-Enforcement*

We conclude this section by observing that strengthening external enforcement implies a higher welfare level in contractual equilibrium. External enforcement becomes stronger if, for instance, the partition P in each stage game becomes finer

exist if Γ were formed by starting with a finite number of finite stage games and then augmenting them with arbitrary contingent transfers, but such speculation is unfounded. Indeed, the optimal stage game outcome might be unenforceable, yet be "virtually enforceable" via an unbounded sequence of transfers, as we show in Appendix Section B2. One might further speculate that if a finite number of finite stage games were augmented with uniformly bounded contingent transfers, then a CEV collection ought to exist, but a bound on transfers can interfere with Theorem 1 in problematic cases. We do not view the lack of a general existence guarantee as a practical problem, as one can work with a near-supremum level for a variant of the CEV definition (see online Appendix Section C.2).

(allowing c to be conditioned on more information about the outcome) or if the set of enforceable production technologies expands. Recalling that the contractual setting is described by (Γ, c^0, π), we can relate two contractual settings most simply by inclusion, holding fixed c^0 and π: Setting $(\tilde{\Gamma}, c^0, \pi)$ is *stronger* than setting (Γ, c^0, π) if $\Gamma \subset \tilde{\Gamma}$. That is, to get a stronger contractual setting we enlarge the set of stage games (and thus the set of available contracts), so all of the items in the weaker technology are retained.

THEOREM 3: *If contractual setting* $(\tilde{\Gamma}, c^0, \pi)$ *is stronger than* (Γ, c^0, π), *and each setting satisfies the conditions in Theorem* 1 *or Theorem* 2, *then the contractual-equilibrium welfare level is weakly higher under* $(\tilde{\Gamma}, c^0, \pi)$.[25]

PROOF:

The result follows from the observation that, in optimization problems (11) and (12), the constraint set under (Γ, c^0, π) is a subset of the constraint set under $(\tilde{\Gamma}, c^0, \pi)$. ■

This conclusion contrasts with some of the prior literature in relational contracts, which has found that under specific assumptions on equilibrium selection, improving external enforcement can reduce welfare. 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 external spot contract. In contrast, contractual equilibrium posits that the parties can always renegotiate both the external contract and their self-enforced arrangements. Thus, when they successfully renegotiate following any history, they agree to an optimal combination of externally enforced and self-enforced elements.

Theorem 3 is in line with empirical studies that find complementarity between the strength of external enforcement and the efficacy of self-enforcement. For example, Johnson, McMillan, and Woodruff (2002) uses the transition of formerly planned economies in Eastern Europe and the Soviet Union, where bureaucratic controls were replaced by more market-oriented legal systems, to examine interactions between the courts and relational contracting. The paper finds that informal arrangements (self-enforcement) are the main basis for contracting by firms in the dataset, and that improvements in legal institutions (enabling better external enforcement) are associated with more effective relational contracting and higher overall productivity. Further, recent studies of inter-firm contracting in developed economies, including Beuve and Saussier (2012), Poppo and Zenger (2002), and Ryall and Sampson (2009), report a positive relation between the extent of "formal contracting" (complexity of the contract and its use of external enforcement) and self-enforcement. In this context, Theorem 3 is directly relevant where empirical variation entails improvements in production technology and monitoring.[26]

[25] A more general version of this result appears in online Appendix Section C.2.

[26] A more general theoretical connection between the use of external enforcement and self-enforcement would require measures of degree in both of these categories and including in the model elements that influence

IV. Option Contracts and the Allocation of Decision Rights

This section continues our analysis of a manager and a worker who can write long-term contracts governing their monitoring technology, introduced in Section IIB. Here we enrich the contractual setting, allowing the parties to construct a menu of options for one of them to verifiably select from, where each option specifies a monitoring level and an externally enforced monetary transfer. As in Section IIB, the monitoring signal is not verifiable, although both the manager and the worker observe it. In this environment with "option contracts" we demonstrate the full power of Theorem 1,[27] and provide insight regarding the optimal allocation of decision rights. Specifically, we find that decision rights are optimally allocated to the manager when the manager has high bargaining power, but to the worker when the worker has high bargaining power. In both cases welfare is maximized by a semi-stationary contract in which the stationary part offers two menu options, while the initial part offers one menu option. Welfare is also higher than in the setting without options, illustrating the complementarity between self-enforcement and external enforcement.

The contracting environment now provides an array of stage games, in which first one party chooses from a menu of two monitoring/payment pairs, (μ^1,p^1) and (μ^2,p^2); then the transfer p^j is made from the manager to the worker; and finally the worker selects effort $a_1 \in \{0,1\}$ under the chosen monitoring technology μ^j.[28] The 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, $((\mu^1,p^1),(\mu^2,p^2)) \in ([0,1] \times \mathbb{R})^2$.

Recall that if the worker exerts low effort then the monitoring signal is low with probability μ, but if the worker exerts high effort then the signal is high for sure; the worker's cost of high effort is β; and the manager incurs monitoring cost $k(\mu)$. Given μ and span d, the worker can be induced to exert high effort if $\delta\mu d \geq (1-\delta)\beta$.

A. *Allocating Decision Rights to the Manager*

We first consider the case in which decision rights are allocated to the manager. The manager can use his discretion to treat the worker differently under disagreement when she is to be rewarded versus when she is to be punished. By Theorem 1, the optimal contract is semi-stationary, with the same menu $((\mu^1,p^1),(\mu^2,p^2))$ specified for every future period. For the current period, parties set a specific monitoring level μ and payment p (that is, a menu $((\mu,p),(\mu,p))$), to maximize their attainable joint value. Theorem 1 instructs us to compute the span by finding the largest fixed point d^* of Λ (optimization problem (11)), and to compute the level

degree, such as contracting costs. Lazzarini, Miller, and Zenger (2004) reports evidence of complementarity in an experiment that examines variations in contracting cost and the length of relationships.

[27] The example in Section IID also applied Theorem 1, but there the monitoring signal was verifiable, so in the optimal contract the worker's incentives came entirely from contractual monetary bonuses. Here the monitoring signal is not verifiable, so the worker must be motivated by relational incentives.

[28] Since the stage game is not simultaneous, technically we must expand the notion of a stage game to allow for simple dynamics, and strengthen Definition 1 to require that an action profile is enforced if it is a subgame perfect equilibrium of the relevant induced game, rather than merely a Nash equilibrium. Since this is intuitive, we do not provide the strengthened formal definitions.

by solving $\Xi(d^*)$ (optimization problem (12)). We focus on the case in which high effort can be implemented, i.e., $d^* \geq ((1-\delta)/\delta)\beta$.

We begin by identifying the stationary part of the optimal contract by finding the largest fixed point of Λ. In optimization problem $\Lambda(d)$, the objective to be maximized is the difference in normalized values $\omega_1(\gamma, a^2, y^2) - \omega_1(\gamma, a^1, y^1)$, by choice of the game γ, action profiles a^1 and a^2, and normalized continuation value mappings y^1 and y^2, subject to enforcement constraints. The constraints ensure that the manager's selection of options and the worker's choice of efforts are incentive compatible. The largest fixed point of Λ will be the span of the optimal contract. For this example, this is a straightforward problem, so here we just state the result (see online Appendix Section C.6 for details).

PROPOSITION 1: *In the setting with options contracts selected by the manager, optimization problem* (11) *has a solution for all* $d \geq 0$. *If* $\delta \geq \beta$ *then* Λ *has a largest fixed point* d^* *satisfying* $d^* \geq ((1-\delta)/\delta)\beta$, *and given by the largest solution to*

$$d = 1 - \beta + \pi_2\left(k(1) - k\left(\frac{(1-\delta)\beta}{\delta d}\right)\right). \tag{13}$$

This span is attained using a stage game with menu items featuring monitoring levels $\mu^1 = 1$ *and* $\mu^2 = (1-\delta)\beta/\delta d^* \leq 1$, *and payments* p^1 *and* p^2 *that satisfy*

$$-p^1 - k(1) = 1 - p^2 - k(\mu^2); \tag{14}$$

and by directing the worker to exert high effort $(a_1^1 = a_1^2 = 1)$ *if the manager selects the correct option* $(a_2^j = (\mu^j, p^j), j = 1,2)$, *but low effort otherwise.*

The stationary part $(\underline{c})$ of the optimal contract specifies a stage game with menu items $(\mu^1, p^1) = (1, p^1)$ and $(\mu^2, p^2) = \left((1-\delta)\beta/\delta d^*, p^1 + 1 + k(1) - k(\mu^2)\right)$. The worker is induced to exert effort both to generate her least favorable and her most favorable payoff under disagreement. The worker's least favored payoff is effectuated partly via a low contractual payment $(p^1 < p^2)$, and partly via a strict monitoring level $(\mu^1 = 1)$ that prevents the worker from earning any rents from effort. The difference between the payments in the two options is made as large as possible, subject to the manager being willing to select Option 2 when the worker is to be rewarded. In fact, equation (14) is the manager's binding incentive constraint for choosing Option 2, reflecting that if the manager deviates, then the worker exerts low effort and the parties coordinate on $z^2(\underline{c})$ for any outcome. Moreover, to maximize p^2, the monitoring level under Option 2 is set to the minimal level μ^2 that induces the worker to exert effort, given the span d^*.

To explain the expression for d (equation (13)), note that under Option 2 the monitoring level μ^2 is already minimized, so the arrangement under Option 2 is already efficient. Without any welfare improvement to negotiate over, the payoffs under agreement and disagreement are the same $(\underline{w}^2 = z^2(\underline{c}))$. The worker's payoff from the action phase when exerting high effort is $\underline{w}_1^2 = (1-\delta)(p^2-\beta) + \delta z_1^2(\underline{c})$, hence $z_1^2(\underline{c}) = p^2 - \beta$.

Under Option 1 the cost of monitoring is maximized, so there is a welfare improvement $k(1) - k(\mu^2)$ to be gained by negotiating to the efficient monitoring level μ^2. The worker's payoff from the action phase when exerting high effort under disagreement is $\underline{w}_1^1 = (1-\delta)(p^1-\beta) + \delta(z_1^1(\underline{c}) + ((1-\delta)/\delta)\beta)$, and she gets her share of the negotiation surplus, so her payoff after negotiating to an agreement from Option 1 is $z_1^1(\underline{c}) = p^1 + \pi_1(k(1) - k(\mu^2))$. The span is thus $d = p^2 - p^1 - \beta - \pi_1(k(1) - k(\mu^2))$. Substituting for $p^2 - p^1$ from equation (14) yields equation (13).

With the span d^* in hand, now we identify the initial part of the optimal contract by solving optimization problem $\Xi(d^*)$. The objective is to maximize the welfare level by choice of the game γ, action profile a, and normalized continuation value mapping y, subject to incentive compatibility constraints. Since we are focusing on the case in which $d^* \geq ((1-\delta)/\delta)\beta$, the solution to $\Xi(d^*)$ is straightforward: the worker should exert high effort, and the cost of the monitoring technology should be minimized subject to the worker's incentive constraint. This entails monitoring level $\mu^* = \mu^2$ and thus optimal welfare level

$$L^* = 1 - \beta - k\left(\frac{(1-\delta)\beta}{\delta d^*}\right).$$

There is no need for two distinct menu items, since the payment cannot be conditioned on the monitoring outcome, so the optimal contract c^* should specify a menu of the form $((\mu^*,p),(\mu^*,p))$. The specific contractual payment p is of no importance, as the players can use voluntary transfers to obtain any desired split of the joint value L^* between them.[29] One possibility is $p = p^2$, which means that contract c^* would specify the same terms as Option 2 for the current period, and thus entail a temporary suspension of Option 1.

In fact, in this setting the optimal contract can be implemented by a stationary contract with the two options $((\mu^1,p^1),(\mu^2,p^2))$ specified for every period. This is possible since Option 2 implements the optimal welfare level L^* under disagreement. By choosing the payment p^2 such that the payoff $z^2(\underline{c})$ accords with their desired division of the surplus L^* (and adjusting p^1 to maintain the optimal span), the parties can achieve their optimal agreement outcome by agreeing each period to implement Option 2, and do this in the same way as Option 2 is implemented under disagreement. In this setting a stationary contract with two options is thus sufficient, where one option has mild monitoring and is selected every period in equilibrium; and the other option has strict monitoring and is selected only if disagreement arises after a low monitoring signal.

The contractual equilibrium is illustrated in Figure 3. The manager's decision rights enable the parties to support a larger span than the contract from Section IIB, where the contractual setting did not allow for options. The span here is at least

[29] The foregoing analysis leaves both p^1 and p as free parameters. As noted, p does not matter at all. Somewhat in contrast, p^1 determines where the value set of the optimal contract is located, although it does not affect its level or span. Nonetheless, at the time the parties agree on their contract, they can use their voluntary transfer during negotiation to offset any change in p^1, if for any reason they select a p^1 that generates a value set that does not contain their desired continuation value.

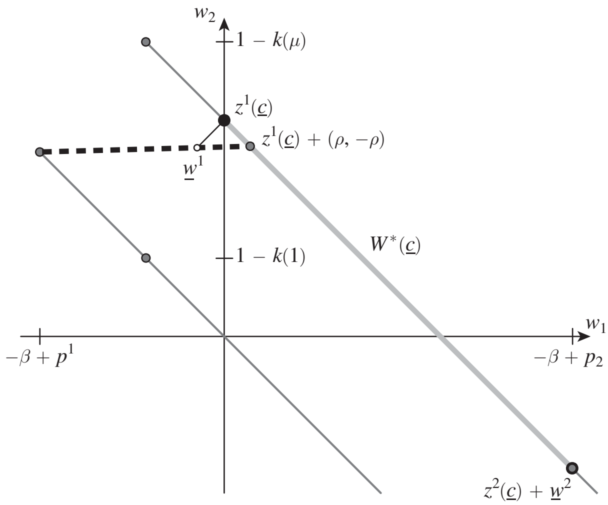

Figure 3. Contractual Equilibrium with Managerial Decision Rights

Notes: Figures 3 and 4 are drawn to scale using the same parameters as in Figure 1. Transfers not pinned down in the analysis are chosen so that $z_1^1(\underline{c}) = 0$. Disagreement point $\underline{w}^1$ is attained by choosing option $(1, p^1)$, playing $a_1 = 1$, and continuing with promised utility $z^1(\underline{c}) + (\rho, -\rho)$ if the signal is $x = 1$, but $z^1(\underline{c})$ if $x = 0$. When $\underline{w}^1$ is the disagreement point, the parties renegotiate to $z^1(\underline{c})$. Disagreement point $\underline{w}^2$ is attained by choosing (μ^*, p^2), playing $a_1 = 1$, and continuing with $z^2(\underline{c})$ if $x = 1$; but with $z^1(\underline{c})$ if $x = 0$. Since $\underline{w}^2$ is efficient, $z^2(\underline{c}) = \underline{w}^2$.

$1 - \beta$, and even greater as the manager's bargaining power π_2 increases, whereas the span without options was merely $\pi_1(1 - \beta)$. The larger span gives the worker higher-powered incentives, which the parties use to reduce their monitoring costs on the equilibrium path. When the manager has higher bargaining power, he takes a greater share of the surplus when renegotiating out of a situation (under disagreement after a low monitoring signal) in which Option 1 is to be chosen, which shifts endpoint $z^1(\underline{c})$ toward a lower worker payoff and enlarges the span.

In practical terms, we can interpret p^1 as the worker's base salary; then she earns a small bonus after low monitoring signals (awarded during renegotiations in return for agreeing to reduce the monitoring level) and a large bonus after high monitoring signals. Only if a disagreement arises after a low monitoring signal does the worker earn merely p^1. While the manager has decision rights, the menu of options constrains him to award either a zero bonus or a large bonus $(p^2 - p^1)$ under disagreement. The large difference between these bonuses is a major contributor to the large span of the optimal contract; the difference is constrained only by the manager's incentive constraint for choosing the right option. This incentive constraint is relatively mild because the zero bonus is paired with high monitoring costs, and because the worker will shirk if the manager chooses the wrong option.

B. *Allocating Decision Rights to the Worker*

In a contractual setting that allows for options contracts, it may be possible to allocate decision rights to either party. In this section we show that it can be optimal to allocate decision rights to the worker, if she has sufficient bargaining power.

The optimization problem $\Lambda(d)$ with worker decision rights has the same objective function and worker's effort incentive constraints as the case with manager decision rights. In place of the manager's incentive constraint, the worker now has an additional incentive constraint for choosing the appropriate option from the menu. In online Appendix Section C.6 we show the following.

PROPOSITION 2: *In the setting with options contracts selected by the worker, optimization problem* (11) *has a solution for all* $d \geq 0$, *and for* $\delta \geq \beta/(\pi_1(1-\beta)+\beta)$ *there is a largest fixed point* d^W *of* Λ *satisfying* $d^W \geq ((1-\delta)/\delta)\beta$. *It is the largest solution to*

$$d = \pi_1\left(1 - \beta + k(1) - k\left(\frac{(1-\delta)\beta}{\delta d}\right)\right). \tag{15}$$

This span is attained using a stage game with menu items featuring monitoring levels $\mu^1 = (1-\delta)\beta/\delta d^W \leq 1$ *and* $\mu^2 = 1$; *payments* p^1 *and* p^2 *that satisfy*

$$p^2 = p^1 + \beta/\mu^1 - \beta; \tag{16}$$

and by directing the worker to exert high effort $(a_1^1 = 1)$ *if Option 1 is correctly selected, but low effort* $(a_1^2 = 0)$ *if Option 2 is correctly selected or if the wrong option is selected.*

The stationary part ($\underline{c}$) of an optimal contract thus specifies a stage game with menu items $(\mu^1,p^1) = ((1-\delta)\beta/\delta d^W, p^1)$ and $(\mu^2,p^2) = (1,p^1+\beta/\mu^1-\beta)$.

Under contract $\underline{c}$ the worker is punished under disagreement by being induced to select Option 1, earn payment p^1, exert high effort, and receive continuation value $z^2(\underline{c})$ if the monitoring signal is high. Since μ^1 is the lowest monitoring level that induces high effort, welfare is maximal (conditional on the span), leaving no welfare improvement to negotiate over. The worker's payoff under both disagreement and agreement is thus $z_1^1(\underline{c}) = p^1 + \beta/\mu^1 - \beta$, where the latter two terms constitute the worker's rent.

In contrast, to reward the worker under contract $\underline{c}$, in disagreement she is induced to select Option 2, earn payment p^2, exert low effort, and receive continuation value $z^2(\underline{c})$. The payment p^2 is set as large as possible relative to p^1, subject to the constraint that the worker is willing to select Option 1 when appropriate; equation (16) expresses this constraint in binding form.[30] In addition, in this case the cost of monitoring is maximized $(\mu^2 = 1)$ in order to punish the manager.[31] This yields a large welfare improvement to be shared when the parties negotiate. The worker gets her share of this improvement, and thus, gets payoff

[30] The worker is deterred from selecting the wrong Option 2 by the threat that the parties will then coordinate on her worst payoff $z^1(\underline{c})$ for any signal; this just suffices when payments satisfy equation (16).

[31] We assume that a high monitoring level μ^2 can be enforced, even if it is intended that the worker should shirk. This is in effect a way for the parties to "burn money" in this setting. An alternative interpretation could be that a third party, e.g., a supplier of monitoring equipment, is entitled to a payment $k(\mu^2)$ under this option, irrespective of whether the equipment is installed.

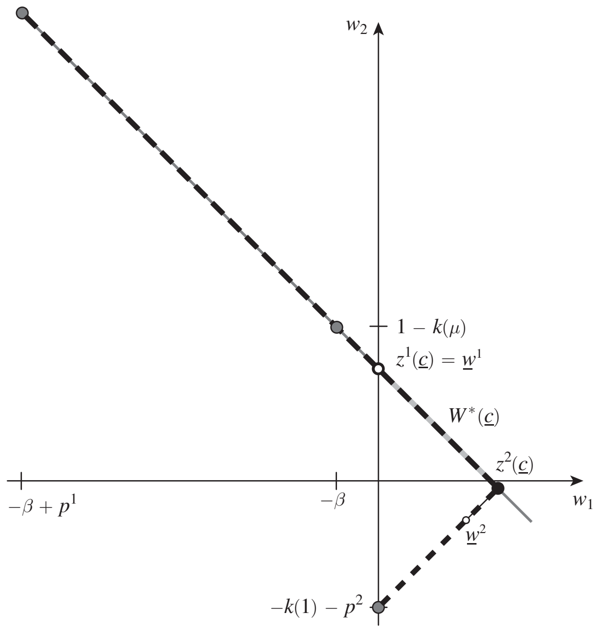

Figure 4. Contractual Equilibrium with Worker Decision Rights

Notes: Disagreement point $\underline{w}^1$ or any other point in W^* is attained by choosing option (μ^1, p^1), playing $a_1 = 1$, and continuing with promised utility $z^2(\underline{c})$ if the signal is $x = 1$, but with $z^1(\underline{c})$ if $x = 0$. Since $\underline{w}^1$ is efficient, $z^1(\underline{c}) = \underline{w}^1$. Disagreement point $\underline{w}^2$ is attained by choosing $(1, p^2)$, playing $a_1 = 0$, and continuing with $z^2(\underline{c})$ regardless of x. When $\underline{w}^2$ is the disagreement point, the parties renegotiate to $z^2(\underline{c})$.

$z_1^2(\underline{c}) = p^2 + \pi_1(1 - \beta - k(\mu^2) + k(1))$. Accounting for equation (16), we then see that the span must satisfy equation (15).

The contractual equilibrium is illustrated in Figure 4. When the worker has a lot of bargaining power, she takes a large share of the surplus when renegotiating out of a situation (under disagreement after a high monitoring signal) in which Option 2 is to be chosen, which shifts endpoint z^2 toward a higher worker payoff and enlarges the span. In contrast, there is no surplus to be negotiated over when Option 1 is to be chosen, so the worker's bargaining power has no effect on endpoint z^1.

Comparing equations (13) and (15), we see that $d^W > d^*$ if π_1 is sufficiently large. Thus, there will be a threshold $\pi_1^* \in (0,1)$ such that allocating decision rights to the worker generates a larger span than allocating decision rights to the manager whenever the worker's bargaining power is sufficiently high $(\pi_1 > \pi_1^*)$. The larger span yields higher-powered incentives, enabling reduced monitoring costs and greater welfare.

Our analysis provides a new explanation for why a contract may optimally allocate limited decision rights to the worker. When the worker has high bargaining power, she will capture a large share of any renegotiation surplus, making it desirable to specify a contract in which the renegotiation surplus is large when the worker is to be rewarded but small when she is to be punished. As we have shown, when decision rights are contractible, allocating them to the worker facilitates such a contract.

V. 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).[32] Levin (2003) showed that with transfers, stationary contracts are optimal in time-invariant environments. Levin also observed that optimal stationary contracts are "strongly optimal" in the sense that the continuation contract at any feasible history is optimal from that point onward. Goldlücke and Kranz (2013) showed that with transfers, perfect monitoring, and no external enforcement, 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 say 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 (2018) also models renegotiation explicitly and allows for disagreements in a repeated game setting, 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.

The literature has shown that optimal relational contracts in time-invariant environments with limited external enforcement may be nonstationary due to one party's limited commitment to a long-term contract (Ray 2002), limited liability (Fong and Li 2017), or persistent private information (Martimort, Semenov, and Stole 2017). No such features are present in the model analyzed here; rather we show that limited external enforcement alone may make the equilibrium contract nonstationary. As noted in the introduction, nonstationarities arise also in the complementary model of Kostadinov (2019).

On the theme of external enforcement operating in concert with self-enforcement, Iossa and Spagnolo (2011) has pointed out that it is common practice to write contracts that contain inefficient clauses, but where these clauses are ignored in equilibrium. They explain this practice by observing that such contracts can be used as a credible threat to sustain a more efficient outcome. Bernheim and Whinston (1998) emphasizes that, when some aspects of performance are unverifiable, it is often optimal to leave other verifiable aspects of performance unspecified, so optimal contracts are less complete than they could have been.[33] In a contractual equilibrium, the optimal contract may entail such flexibility. Further, flexibility in the form of options can be valuable, and then the allocation of decision rights is relevant.

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

[33] Iossa and Spagnolo (2011) examines a repeated principal-agent model in which, in each period, players have the option to trigger penalties specified by the contract. Long-term 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) examines a class of two-period contracting problems with both external enforcement and self-enforcement.

Baker, Gibbons, and Murphy (2011) also demonstrates how allocation of such rights matters in relational contracting, but via a channel 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 et al. (2019) analyzes self-enforced agreements that facilitate efficient 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 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.[34] Researchers have shown that the hold-up problem can be alleviated in some short-term trading relationships, in particular in settings of "own-investment" (e.g., Aghion, Dewatripont, and Rey 1994; Nöldeke and Schmidt 1995; Edlin and Reichelstein 1996). Results in this literature rely on complementarities, specifically that investment decisions influence the value of trade. In our model, as with most in the relational-contracting literature, all actions that affect the surplus occur at the same time, 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 actions taken in one period influence the payoffs received in future periods.

VI. Conclusion

This paper makes four related contributions. First, we introduce a flexible model of long-term contractual relationships with external enforcement. The contracting parties can write an arbitrary nonstationary long-term contract that specifies a stage game for them to play as a function of the verifiable history. The details of the contracting environment are represented by the collection of available stage games. We extend *contractual equilibrium* (Miller and Watson 2013) to this environment, to allow for renegotiation, bargaining power, and the possibility of disagreement.

Second, we show that *semi-stationary* contracts are optimal in two important classes of contracting environments: those with *no verifiable information*, and those with *externally enforced contingent transfers*. In a semi-stationary contract, there are special terms for the present period, conducive to high payoffs; and there are stationary terms for all future periods, inducing the greatest span of continuation val-

[34] Prominent entries include Hart and Moore (1988, 1999), Nöldeke 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.

ues consistent with incentives. Unlike arbitrary nonstationary long-term contracts, semi-stationary contracts are tractable, and we provide a method for optimizing them.

Third, we show that, in contractual equilibrium, self-enforcement and external enforcement are always complementary: if the external enforcement becomes stronger, the welfare level in contractual equilibrium becomes higher.

Finally, we analyze a principal-agent model with moral hazard, where the manager and the worker can contractually specify their monitoring technology. In the simplest case, with no verifiable information, we show that the optimal contract specifies mild monitoring for the current period and intense monitoring for all future periods. In each period, the parties renegotiate back to this same contract, so on the equilibrium path they always operate under mild monitoring. The intense monitoring specified for the future facilitates incentives for the worker. We analyze several extensions of this model, most notably by allowing the parties' contract to allocate decision rights over the monitoring technology. The fact that the decision (one party's choice from a menu of monitoring and payment combinations) is verifiable enhances the power of incentives. Depending on their relative bargaining power, it can be optimal to allocate decision rights to either the manager or the worker.

We hope these contributions have laid the groundwork for continued research on long-term contracts and the interaction between external enforcement and self-enforcement. While our results on semi-stationarity may apply to many interesting cases, there are many others which may require more complicated nonstationary contracts, for instance, if there are limited-liability constraints or if the technological environment itself is nonstationary.

Appendix A. Proof of the Main Result

This section proves Theorem 1. The proof proceeds with a series of lemmas, interspersed with some guiding comments and statements about notation.

LEMMA 2: Λ *has a maximal fixed point*, *denoted* d^*.

PROOF OF LEMMA 2:

Recall that $\omega(\gamma, \alpha, y)$ denotes the normalized continuation value if in the current period α is played in stage game $\gamma = (A, X, \lambda, u, P)$ and the continuation value in the next period is given by $y : X \to \mathbb{R}^2_0$. This was defined in Section IIIA. For a given span d, let γ, y^1, y^2, α^1, and α^2 solve optimization problem (11) to determine $\Lambda(d)$, and let u be the payoff function for stage game γ. From the definition of ω we have that

$$\Lambda(d) = (1-\delta)\,u_1(\alpha^2) + \delta\bar{y}^2_1(\alpha^2) - \pi_1(1-\delta)\big(u_1(\alpha^2) + u_2(\alpha^2)\big)$$
$$- \big[(1-\delta)\,u_1(\alpha^1) + \delta\bar{y}^1_1(\alpha^1)\big] + \pi_1(1-\delta)\big(u_1(\alpha^1) + u_2(\alpha^1)\big).$$

Recall that we assumed joint payoffs in stage games are bounded uniformly below by $-\vartheta$ and above by ϑ. Therefore $u_1(\alpha^2) + u_2(\alpha^2) \geq -\vartheta$ and $u_1(\alpha^1) + u_2(\alpha^1) \leq \vartheta$, and we have

$$\text{(A1)} \qquad \Lambda(d) \leq (1-\delta)\,u_1(\alpha^2) + \delta\bar{y}_1^2(\alpha^2) - \left[(1-\delta)\,u_1(\alpha^1) + \delta\bar{y}_1^1(\alpha^1)\right] + 2\pi_1(1-\delta)\vartheta.$$

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 payoff exceeds $-\vartheta$, enforcement of α^2 (in particular that player 2 cannot gain by deviating to α_2^1), and that the joint stage-game payoff is no greater than ϑ:

$$-\left[(1-\delta)\,u_1(\alpha^1) + \delta\bar{y}_1^1(\alpha^1)\right] \leq -(1-\delta)\,u_1(\alpha_1^2,\alpha_2^1) - \delta\bar{y}_1^1(\alpha_1^2,\alpha_2^1),$$

$$0 \leq (1-\delta)\,u_1(\alpha_1^2,\alpha_2^1) + (1-\delta)\,u_2(\alpha_1^2,\alpha_2^1) + (1-\delta)\vartheta,$$

$$0 \leq (1-\delta)\,u_2(\alpha^2) + \delta\bar{y}_2^2(\alpha^2) - (1-\delta)u_2(\alpha_1^2,\alpha_2^1) - \delta\bar{y}_2^2(\alpha_1^2,\alpha_2^1),$$

$$0 \leq -(1-\delta)\,u_2(\alpha^2) - (1-\delta)\,u_1(\alpha^2) + (1-\delta)\vartheta.$$

Summing these inequalities yields

$$-\left[(1-\delta)\,u_1(\alpha^1) + \delta\bar{y}_1^1(\alpha^1)\right] \leq -\delta\bar{y}_1^1(\alpha_1^2,\alpha_2^1) - \delta\bar{y}_2^2(\alpha_1^2,\alpha_2^1) + \delta\bar{y}_2^2(\alpha^2) - (1-\delta)\,u_1(\alpha^2) + 2(1-\delta)\vartheta.$$

Substituting the bracketed left-side terms into equation (A1) and simplifying, we obtain

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

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

$$\text{(A2)} \qquad \Lambda(d) \leq 2(1+\pi_1)(1-\delta)\vartheta + \delta d.$$

In words, $\Lambda(d)$ is bounded above by a line with slope $\delta < 1$. We thus know that $\Lambda(d) < d$ for all $d > \bar{d}$ where $\bar{d}$ solves $\bar{d} = 2(1+\pi_1)(1-\delta)\vartheta + \delta\bar{d}$. Clearly Λ is increasing, and since $\Lambda(0) \geq 0$ we know that the restriction of Λ to subdomain $[0,\bar{d}]$ maps to the same set, and thus Λ has a maximal fixed point d^* by Tarski's fixed-point theorem. ∎

As noted in Section IIIA, let $\underline{\gamma} = (\underline{A}, \underline{X}, \underline{\lambda}, \underline{u}, \underline{P})$, y^1, y^2, α^1, and α^2 denote any solution to optimization problem (11) for Λ evaluated at d^*. Let $\gamma^* = (A^*, X^*, \lambda^*, u^*, P^*)$, y^*, and α^* denote any solution to optimization problem (12) for Ξ evaluated at d^*. Define $\underline{c}$ to be the stationary external contract that specifies stage game $\underline{\gamma}$ in every period, and define c^* to be the semi-stationary contract that specifies stage game γ^* for the current period and then transitions to $\underline{c}$. We will eventually demonstrate that c^* is optimal.

For any stage game $\gamma = (A, X, \lambda, u, P)$, let $Z(\gamma, d)$ denote the set of normalized continuation values that can be achieved in the induced game where players engage in γ and then coordinate on continuation values in the set $\mathbb{R}^2_0(d)$, for a given span d:

$$Z(\gamma, d) \equiv \left\{\omega(\gamma, \alpha, y) \,|\, y : X \to \mathbb{R}^2_0(d); \alpha \text{ is enforced relative to } \gamma \text{ and } y\right\}.$$

By definition of $\underline{\gamma}$, we have $\operatorname{span}(Z(\underline{\gamma}, d^*)) = d^*$ and $Z(\underline{\gamma}, d^*)$ attains its span.

LEMMA 3: *For any $L \in \mathbb{R}$, take as given a collection $\mathcal{W} = \{W(c')\}_{c' \in C}$ with at least one nonempty set and satisfying $w_1 + w_2 = L$ for every $w \in W(c')$ and $c' \in C$. Let d be any number satisfying $d \geq \sup\{\operatorname{span}(W(c')) \,|\, c' \in C\}$. Consider any $c \in C$ and $w \in \mathbb{R}^2$ such that w is c-supported relative to $\mathcal{W}$. It is the case that $w_1 + w_2 \leq (1-\delta)\Xi(d) + \delta L$.*

PROOF OF LEMMA 3:

Let $(A, X, \lambda, u, P) = g(c)$. From the definition of c-support, there exists $\alpha \in \Delta A$ and $y : X \to \mathbb{R}^2$ such that $y(x) \in \operatorname{co} W(c|x)$ for all $x \in X$, α is enforced relative to $g(c)$ and y, and $w = (1-\delta)u(\alpha) + \delta \bar{y}(\alpha)$. Because $d \geq \operatorname{span}(W(c|x))$, every point in $W(c|x)$ has joint value L for all $x \in X$, and $c|\cdot$ is P-measurable, we can find a P-measurable function $b : X \to \mathbb{R}^2_0$ such that

$$\text{(A3)} \qquad \operatorname{co} W(c|x) \subset \mathbb{R}^2_0(d) + \frac{1-\delta}{\delta} b(x) + \pi L$$

for every $x \in X$. The corresponding expected transfer function $\bar{b} : A \to \mathbb{R}^2_0$ is given by $\bar{b}(a) \equiv E_x[b(x) | x \sim \lambda(a)]$ for every $a \in A$. Let $\gamma' \equiv (A, X, \lambda, u + \bar{b}, P)$. Because stage game γ' merely adds P-measurable transfers to stage game γ, we know $\gamma' \in \Gamma$ by the assumption of externally enforced contingent transfers.

Let us define $y' : X \to \mathbb{R}^2$ by

$$\text{(A4)} \qquad y'(x) \equiv y(x) - \pi L(\mathcal{W}) - \frac{1-\delta}{\delta} b(x)$$

for every $x \in X$. Expressions (A3) and (A4) imply that y' is a function from X to $\mathbb{R}^2_0(d)$. Substituting for y', we see that the induced game

$$\text{(A5)} \qquad \left\langle A, (1-\delta)u(\cdot) + (1-\delta)\bar{b}(\cdot) + \delta \bar{y}'(\cdot)\right\rangle$$

is equivalent to induced game $\langle A, (1-\delta)u(\cdot) + \delta \bar{y}(\cdot)\rangle$ up to the constant $\pi\delta L(\mathcal{W})$ in the payoff function, which establishes that α is enforced relative to γ' and y'. Because γ' and y' are feasible in optimization problem (12) for $\Xi(d)$, we conclude that $u_1(\alpha) + \bar{b}_1(\alpha) + u_2(\alpha) + \bar{b}_2(\alpha) \leq \Xi(d)$. Since $\bar{b}_1(\alpha) + \bar{b}_2(\alpha) = 0$, this means $u_1(\alpha) + u_2(\alpha) \leq \Xi(d)$. Recalling that $w = (1-\delta)u(\alpha) + \delta\bar{y}(\alpha)$, we have thus established $w_1 + w_2 \leq (1-\delta)\Xi(d) + \delta L$. ∎

Hereafter, it is useful to represent bargaining self-generation with the operator

$$B(\hat{c},\mathcal{W}) \equiv \left\{\underline{w} + \pi\left(L(\mathcal{W}) - \underline{w}_1 - \underline{w}_2\right) \middle| \underline{w} \text{ is } \hat{c}\text{-supported relative to } \mathcal{W}\right\}, \tag{A6}$$

assuming $L(\mathcal{W})$ exists and there exists a $\hat{c}$-supported value; otherwise, let $B(\hat{c},\mathcal{W}) \equiv \varnothing$. Then a collection $\mathcal{W}$ is BSG if $W(c) \subset B(c,\mathcal{W})$ for every $c \in C$. The next lemma identifies the collection described at the end of Section IIIA.

LEMMA 4: *There is a BSG collection* $\mathcal{W} = \{W(c)\}_{c\in C}$ *for which* $\text{span}(W(\underline{c})) = d^*$, $W(\underline{c})$ *attains its span*, $L(\mathcal{W}) = \Xi(d^*)$, *and* $W(c^0) \neq \varnothing$.

PROOF OF LEMMA 4:

First, any action profile α that is enforced relative to $\underline{\gamma}$ and some $y : \underline{X} \to \mathbb{R}^2_0(d^*)$ is also enforced relative to $\underline{\gamma}$ and $y(\cdot) + (k, \Xi(d^*) - k)$, for any constant $k \in \mathbb{R}$, because the two induced games are equivalent up to a constant in the payoffs. Suppose that

$$W(\underline{c}) = \left(k, \Xi(d^*) - k\right) + \left\{(0,0),(d^*,-d^*)\right\} \tag{A7}$$

and let us presume for now that $L(\mathcal{W}) = \Xi(d^*)$. By writing the resulting payoff in the induced game for any enforced α and comparing the definitions of operators B and Z, a little algebra reveals that $z \in Z(\underline{\gamma}, d^*)$ if and only if $z + \Xi(d^*)\pi + \delta(k - \Xi(d^*)\pi_1)(1,-1) \in B(\underline{c},\mathcal{W})$. In other words,

$$B(\underline{c},\mathcal{W}) = Z(\underline{\gamma}, d^*) + \Xi(d^*)\pi + \delta\left(k - \Xi(d^*)\pi_1\right)(1,-1). \tag{A8}$$

Since $B(\underline{c},\mathcal{W})$ is a translation of $Z(\underline{\gamma}, d^*)$, it attains its span d^*. We have presumed that the level of $\mathcal{W}$ is $\Xi(d^*)$, so $w_1 + w_2 = \Xi(d^*)$ for every $w \in B(\underline{c},\mathcal{W})$. Therefore, the endpoints of $B(\underline{c},\mathcal{W})$ can be written as a set

$$\left(k', \Xi(d^*) - k'\right) + \left\{(0,0),(d^*,-d^*)\right\} \tag{A9}$$

for some $k' \in \mathbb{R}$. An implication is that equation (A8) implicitly defines a mapping from k to k' (compare expressions (A7) and (A9)). Clearly it is a contraction mapping and its fixed point is $k^* = \Xi(d^*)\pi_1 + \omega_1(\underline{\gamma}, \alpha^1, y^1)/(1-\delta)$, where α^1 and y^1 are given by the solution of optimization problem (11) (noting that $\omega_1(\underline{\gamma}, \alpha^1, y^1)$ is the endpoint of $Z(\underline{\gamma}, d^*)$ that favors player 2). Setting

$$W(\underline{c}) \equiv \left(k^*, \Xi(d^*) - k^*\right) + \left\{(0,0),(d^*,-d^*)\right\},$$

we thus have $W(\underline{c}) \subset B(\underline{c},\mathcal{W})$ regardless of how we define $W(c)$ for $c \neq \underline{c}$.

Next we specify $W(c^0)$. Let $\gamma^0 = (A^0, X^0, \lambda^0, u^0, P^0)$ denote the stage game that default contract c^0 specifies for every period, and let α^0 be a Nash equilibrium of this stage game (which we have assumed exists). Let $W(c^0)$ be the singleton set specified as follows:

$$W(c^0) \equiv \left\{u^0(\alpha^0) + \pi\left(\Xi(d^*) - u_1^0(\alpha^0) - u_2^0(\alpha^0)\right)\right\}.$$

It is evident that $W(c^0) \subset B(c^0, \mathcal{W})$ under our presumption that the level of $\mathcal{W}$ is $\Xi(d^*)$.

So far we have specified $W(\underline{c})$ and $W(c^0)$. For every other contract $c \notin \{\underline{c}, c^0\}$, specify $W(c) = \varnothing$, which completes the construction of $\mathcal{W}$. As verified above, the BSG conditions hold, presuming that the level of $\mathcal{W}$ is $\Xi(d^*)$.

Finally, we justify our presumption that $L(\mathcal{W}) = \Xi(d^*)$. Recall that γ^*, y^*, and α^* solve optimization problem (12) for Ξ evaluated at d^*. This means y^* maps to $\mathbb{R}_0^2(d^*)$, α^* is enforced relative to γ^* and y^*, and $\Xi(d^*) = u_1^*(\alpha^*) + u_2^*(\alpha^*)$. Because $\text{span}(W(\underline{c})) = \text{span}(Z(\underline{\gamma}, d^*)) = d^*$, we know that $y^*(x) + (k^*, \Xi(d^*) - k^*) \in W(\underline{c})$ for every $x \in X^*$. Therefore, noting that α^* is enforced relative to γ^* and $y^* + (k^*, \Xi(d^*) - k^*)$, we have that continuation value $w = (1-\delta)\,u^*(\alpha^*) + \delta\bar{y}^*(\alpha^*) + \delta(k^*, \Xi(d^*) - k^*)$ is c^*-supported relative to $\mathcal{W}$. It is clearly the case that $w_1 + w_2 = \Xi(d^*)$.

By the construction of $\mathcal{W}$ we have $\sup\{\text{span}(W(c'))\,|\,c' \in C\} = \text{span}(W(\underline{c})) = d^*$. Letting $L = \Xi(d^*)$ and $d = d^*$, Lemma 3 implies that no contract can support, relative to $\mathcal{W}$, a joint value in excess of $\Xi(d^*)$. Therefore

$$\max\left\{w_1 + w_2\,|\,c \in C \text{ and } w \text{ is } c\text{-supported relative to } \mathcal{W}\right\} = \Xi(d^*).$$

We have thus constructed a BSG collection with the required properties. ∎

The BSG collection constructed in Lemma 4 is our candidate CEV collection. To demonstrate that it is, in fact, a CEV collection, we must show that there is no other BSG collection that has a strictly higher level. We will do this by showing that the maximal span of BSG sets is d^* and then by showing that $\Xi(d^*)$ obtains the maximal joint value. Let

$$\hat{d} \equiv \sup\left\{\text{span}\left(W(c)\right)|\,\mathcal{W} = \left\{W(c')\right\}_{c' \in C} \text{ is a BSG collection and } c \in C\right\}.$$

We will compare $\hat{d}$ to d^*. The next lemma is the key step, where externally enforced contingent transfers are used to limit the range of y to a single set of continuation values.

LEMMA 5: *For every BSG collection $\mathcal{W} = \{W(c')\}_{c' \in C}$ and for every $c \in C$, there exists $\gamma' \in \Gamma$ such that $W(c) \subset Z(\gamma', \hat{d}) + \pi L(\mathcal{W})$.*

PROOF OF LEMMA 5:

Take as given a BSG collection $\mathcal{W}$ and a contract $c \in C$, and let $(A, X, \lambda, u, P) = g(c)$. Because $\hat{d} \geq \text{span}(W(c|x))$, every point in $W(c|x)$ has joint value $L(\mathcal{W})$ for all $x \in X$, and $c|\cdot$ is P-measurable, we can find a P-measurable function $b : X \to \mathbb{R}^2_0$ such that

$$\text{(A10)} \qquad \text{co}\, W(c|x) \subset \mathbb{R}^2_0(\hat{d}) + \frac{1-\delta}{\delta} b(x) + \pi L(\mathcal{W})$$

for every $x \in X$. Let $\gamma' \equiv (A, X, \lambda, u + \bar{b}, P)$. Because stage game γ' merely adds P-measurable transfers to stage game γ, we know that $\gamma' \in \Gamma$ by the assumption of externally enforced contingent transfers.

Consider any $w \in W(c)$. From the BSG condition, there exists a c-supported continuation value $\underline{w}$ such that $w = \underline{w} + \pi(L(\mathcal{W}) - \underline{w}_1 - \underline{w}_2)$. From the definition of c-supported, there exists $\alpha \in \Delta A$ and $y : X \to \mathbb{R}^2$ such that $y(x) \in \text{co}\, W(c|x)$ for all $x \in X$, α is enforced relative to $g(c)$ and y, and $\underline{w} = (1-\delta)u(\alpha) + \delta \bar{y}(\alpha)$. Following steps in the proof of Lemma 3, we define $y' : X \to \mathbb{R}^2$ by equation (A4) and observe that, from this and expression (A10), y' is a function from X to $\mathbb{R}^2_0(\hat{d})$. Substituting for y', we see that induced game (A5) is equivalent to induced game $\langle A, (1-\delta)u(\cdot) + \delta \bar{y}(\cdot)\rangle$ up to the constant $\pi\delta L(\mathcal{W})$ in the payoff function, which establishes that α is enforced relative to γ' and y', and therefore $\omega(\alpha, \gamma', y') \in Z(\gamma', \hat{d})$.

We conclude by comparing w and $\omega(\alpha, \gamma', y')$. From $w = \underline{w} + \pi(L(\mathcal{W}) - \underline{w}_1 - \underline{w}_2)$, substituting for $\underline{w}$ and using the fact that $\bar{y}_1(\alpha) + \bar{y}_2(\alpha) = L(\mathcal{W})$, a little algebra yields

$$w = (1-\delta)\big(\pi_2 u_1(\alpha) - \pi_1 u_2(\alpha)\big)(1,-1) + (1-\delta)\pi L(\mathcal{W}) + \delta \bar{y}(\alpha).$$

Substituting for $\bar{y}$ using the expectation of equation (A4), $\bar{b}_1(\alpha) + \bar{b}_2(\alpha) = 0$, and $\pi_1 + \pi_2 = 1$, we rearrange terms to get

$$w = (1-\delta)\Big(\pi_2\big(u_1(\alpha) + \bar{b}_1(\alpha)\big) - \pi_1\big(u_2(\alpha) + \bar{b}_2(\alpha)\big)\Big)(1,-1)$$
$$+ \delta \bar{y}'(\alpha) + \pi L(\mathcal{W}),$$

which is $\omega(\alpha, \gamma', y') + \pi L(\mathcal{W})$. We have thus established that $w \in Z(\gamma', \hat{d}) + \pi L(\mathcal{W})$. ∎

LEMMA 6: $\hat{d} = d^*$.

PROOF OF LEMMA 6:

Consider any BSG collection $\mathcal{W} = \{W(c')\}_{c' \in C}$ and any contract $c \in C$. From Lemma 5, there exists $\gamma' \in \Gamma$ such that $W(c) \subset Z(\gamma', \hat{d}) + \pi L(\mathcal{W})$, which implies that $\text{span}(W(c)) \leq \text{span}(Z(\gamma', \hat{d}))$. We also know that $\text{span}(Z(\gamma', \hat{d})) \leq \Lambda(\hat{d})$ because Λ optimizes over the stage game in addition to the enforced action profile. Therefore we have $\text{span}(W(c)) \leq \Lambda(\hat{d})$. Because this weak inequality holds for every external contract and every BSG collection, it also holds at the supremum value, so $\hat{d} \leq \Lambda(\hat{d})$. Because Λ is increasing, satisfies $\Lambda(d) < d$ for all $d > \bar{d}$, and its restriction to subdomain $[0, \bar{d}]$ maps to the same set, it must have a fixed point that weakly exceeds $\hat{d}$, implying that $\hat{d} \leq d^*$. From Lemma 4, a BSG collection exists in which the span d^* is attained, so $\hat{d} = d^*$. ∎

LEMMA 7: *Every BSG collection $\mathcal{W}$ has the property that $L(\mathcal{W}) \leq \Xi(d^*)$.*

PROOF OF LEMMA 7:

Suppose to the contrary there is a BSG collection $\mathcal{W}$ such that $L(\mathcal{W}) > \Xi(d^*)$. Then there must exist a contract $c \in C$ and a value w that is c-supported relative to $\mathcal{W}$, such that $w_1 + w_2 = L(\mathcal{W}) > \Xi(d^*)$. From Lemma 6, we have $d^* \geq \sup\{\text{span}(W(c'))|c' \in C\}$. Applying Lemma 3 with $d = d^*$ and $L = L(\mathcal{W})$ then yields $w_1 + w_2 \leq (1-\delta)\Xi(d) + \delta L(\mathcal{W}) < L(\mathcal{W})$, a contradiction. ∎

To complete the proof of Theorem 1, simply combine Lemmas 4 and 7. Lemma 7 implies that the level of the BSG collection $\mathcal{W}$ identified by Lemma 4 is maximal among the set of BSG collections, and therefore $\mathcal{W}$ is a CEV collection. The maximal CEV collection contains all of the continuation values in $\mathcal{W}$, so the semi-stationary contract c^* identified by Lemma 4 is optimal. ∎

Appendix B. Foundations and Technical Notes

This section begins with a description of contractual equilibrium in terms of strategies in a hybrid game in which stage-game actions are modeled noncooperatively and interaction in the negotiation phase is modeled cooperatively. In Section B2 we discuss technical issues regarding existence and properties of equilibrium, and in Section B3 we comment on the connection between the hybrid model and fully noncooperative models.

B1. *Contractual Equilibrium in Terms of Strategies*

Our hybrid model requires a generalized notion of strategy, called a *regime*, specifying both individual actions in the action phase and joint decisions in the negotiation phase, conditional on the public history. We develop conditions for a contractual equilibrium regime that correspond exactly to the conditions for a CEV collection in Section I. Variations and related results are provided in online Appendix Section C.2.

Recall that play in a single period t consists of the negotiated external contract c^t and transfer m^t (equal to $\hat{c}^t$ and zero in disagreement), the action profile a^t, the outcome x^t, and the unverifiable random draw of the randomization device, which we denote ϕ^t. Let $\psi = (c^t m^t x^t \phi^t)_{t=1}^T$ denote the public history of interaction through any given period T. The history to the action phase of a given period t can be expressed as ψcm, where ψ is the history to the end of period $t-1$ (the null history if $t = 1$) and c and m are jointly chosen in the negotiation phase of period t. Likewise, for a T-period history ψ we write $\psi cmx\phi$ as the $T+1$-period history that appends ψ with joint decision c and m, outcome x, and random draw ϕ in period $T+1$. Define $\kappa(\psi)$ to be the external contract inherited in the period following history ψ. That is, for $\psi = \psi' cmx\phi$, we have $\kappa(\psi' cmx\phi) \equiv c|x$, and if ψ is the null history then $\kappa(\psi) = c^0$. Note that in the period following history ψ, disagreement is represented by selection of $c = \kappa(\psi)$ and $m = 0$.[35]

The joint selection of c^t and m^t is given by functions r^{c} and r^{m} of the public history ψ. The mixed action profile is specified by a function r^{a} of the history to the action phase ψcm. Thus a regime is given by $r = (r^{\mathrm{c}}, r^{\mathrm{m}}, r^{\mathrm{a}})$.

For any contract c, let us write $(A(c), X(c), \lambda(\cdot\,; c), u(\cdot\,; c), P(\cdot\,; c)) = g(c)$ so that we can refer to elements of the stage game in reference to c. Given a T-period history ψ, let $v(\psi; r)$ denote the continuation value following ψ, conditional on the players behaving according to r from this point. That is, $v(\psi; r)$ is the expected value of $\sum_{t=T+1}^{\infty} \delta^{t-T-1}(1-\delta)(m^t + u(a^t; c^t))$, with the expectation taken over the infinite history that begins with ψ. Let $v^{\mathrm{a}}(\psi cm, \alpha; r)$ denote the continuation value from the action phase of a period following history ψcm, conditional on action profile α played in the current period and the players behaving according to r from the next period. From these definitions we have

$$v^{\mathrm{a}}(\psi cm, \alpha; r) = (1-\delta)u(\alpha; c) + \delta E_{x,\phi}\big[v(\psi cmx\phi; r) | x \sim \lambda(\alpha; c), \phi \sim U[0,1]\big].$$

Further, define $\underline{v}(\psi; r) = v^{\mathrm{a}}(\psi\kappa(\psi)0, r^{\mathrm{a}}(\psi\kappa(\psi)0); r)$ as the disagreement point for negotiation in the period following ψ.

For any T-period public history ψ, let $r|\psi$ denote the *continuation regime* following ψ; this is a function of the histories from period $T+1$. Finally, let us call a public history ψ *negotiation-consistent with regime* r if for each period in this history, play in the negotiation phase was either as prescribed by r^{c} and r^{m} or it was the disagreement outcome. That is, for any subhistory $\psi' cm$ (a truncation of ψ), it must be that either $c = r^{\mathrm{c}}(\psi')$ and $m = r^{\mathrm{m}}(\psi')$, or $c = \kappa(\psi')$ and $m = 0$. Note that ψ may entail deviations from r^{a} in the action phase. Call a history ψcm to the action phase *negotiation-consistent* if it has the same property.

The conditions described next will be applied to only the subset of histories that are negotiation-consistent with the regime being evaluated. The reason is technical and relates to existence of equilibrium, which we discuss in Section B2.[36] Call a regime r *incentive compatible in the action phase* if for every history ψcm that is negotiation-consistent with r, neither player would gain by unilaterally deviating from r^{a} in the action phase that follows. That is, for each player i and any action $a_i' \in A_i(c)$, it is the case that $v_i^{\mathrm{a}}(\psi cm, r^{\mathrm{a}}(\psi cm); r) \geq v_i^{\mathrm{a}}(\psi cm, (a_i', r_{-i}^{\mathrm{a}}(\psi cm)); r)$.

Because the hybrid model accounts for behavior in the negotiation phase cooperatively, the equilibrium conditions for this phase are expressed in terms of a bargaining solution, namely the generalized Nash solution with fixed bargaining weights $\pi = (\pi_1, \pi_2)$. We assume that the players negotiate over both the external contract and the self-enforced arrangements. Internal consistency captures the idea

[35] It does not matter for our analysis that our accounting of histories does not differentiate between disagreement and an agreement to keep the inherited external contract and make no transfer.

[36] One could use a stronger notion of equilibrium that requires incentive-compatibility and internal bargain-consistency after all histories, not just those that are negotiation-consistent; this would correspond to a stronger version of CEV that requires $W(c) \neq \varnothing$ for all c. Existence would not be assured for as many contractual settings, and the modified CEV conditions would be less convenient to apply, but otherwise the difference is inconsequential for applications.

that the players may consider altering their regime to select any contractual arrangement for the current period that, from the start of the next period, reverts back to specifications of their current regime (continuing as though the history were some other that is negotiation-consistent with this regime).

To be precise, for a given regime r and after any history ψ, the players contemplate choosing any contract c, transfer m, and action profile $\alpha \in \Delta A(c)$, and then continuing from the next period as though in some other regime r'. Call (c, m, α, r') *comparable with r following* ψ if two conditions hold. First, $v_i^{\mathrm{a}}(\psi cm, \alpha; r') \geq v_i^{\mathrm{a}}(\psi cm, (a_i', \alpha_{-i}); r')$ for $a_i' \in A_i(c)$ and $i = 1, 2$, so behavior in the current period is incentive-compatible. Second, for every $x \in X(c)$ and $\phi \in [0, 1]$, there is a history ψ' that is negotiation-consistent with r such that $\kappa(\psi') = c|x$ and $r'|\psi cmx\phi = r|\psi'$. That is, in regime r' after history $\psi cmx\phi$, the parties behave as if they were in regime r after history ψ'.

The bargaining solution requires that r solves the problem of maximizing the joint value over all such comparable arrangements, the bargaining surplus is defined relative to the disagreement point, and the surplus is divided according to the bargaining weights. That is, letting ℓ denote the maximum of $v_1^{\mathrm{a}}(\psi cm, \alpha; r) + v_2^{\mathrm{a}}(\psi cm, \alpha; r)$ over all (c, m, α, r') that are comparable with r following ψ, we require $v(\psi; r) = \underline{v}(\psi; r) + \pi(\ell - \underline{v}_1(\psi; r) - \underline{v}_2(\psi; r))$. Call regime r *internally bargain-consistent* if this condition holds for every ψ that is negotiation-consistent with r. Clearly ℓ is independent of ψ, so every internally bargain-consistent regime has a single value of ℓ which we call the regime's *level*.

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. To relate the CE definition in terms of strategies to the recursive formulation of CEV collections, let us define for any regime r a collection $\mathcal{V}(r) = \{V(c; r)\}_{c \in C}$ by $V(\hat{c}; r) \equiv \{v(\psi; r) | \psi$ is negotiation-consistent with r and $\kappa(\psi) = \hat{c}\}$ for every $\hat{c} \in C$. Online Appendix Section C.1 establishes the following result.

LEMMA 8: *If r is a contractual equilibrium then* $\mathcal{V}(r)$ *is a CEV collection. If* $\mathcal{W}$ *is a CEV collection then there exists a contractual equilibrium regime r satisfying* $V(c; r) \subset W(c)$ *for every* $c \in C$.

B2. *Technical Issues Regarding Existence*

Two technical issues have arisen in our analysis: $W(c) = \varnothing$ is possible for some c in a CEV collection, and it is difficult to find primitive conditions that guarantee existence. We elaborate with two examples.

Consider first a principal-agent setting in which the agent (player 1) must choose effort $a_1 \geq 0$ at increasing cost, effort is verifiable, and contingent transfers are externally enforced. Consider a contract that, for some threshold $\underline{a}_1 > 0$, specifies a bonus if $a_1 > \underline{a}_1$ and no bonus otherwise. For a large enough bonus, this contract puts the agent in the position of having no best response in the effort subgame. This issue arises naturally in many standard contracting and mechanism-design models, where the typical remedy is to disregard such contracts/mechanisms. In our study, such a problematic contract c has $W(c) = \varnothing$ and correspondingly we do not include c in the incentive-compatibility check (also c would not arise as an inherited contract in negotiation-consistent histories of the hybrid model).

We can rule out examples like the one just described by limiting attention to finite stage games, but existence issues remain. The second example features a class of stage games with externally enforced contingent transfers τ and τ', given by the following payoff matrix:

	Left	Right	Out
Up	$0 + \tau', 1 - \tau'$	$1 + \tau, 0 - \tau$	$0 + \tau', 0 - \tau'$
Down	$2 + \tau', 0 - \tau'$	$0 + \tau', 1 - \tau'$	$0 + \tau', 0 - \tau'$
Out	$0 + \tau', 0 - \tau'$	$0 + \tau', 0 - \tau'$	$0 + \tau', 0 - \tau'$

The partition P is as illustrated by the cell boundaries: the enforcer can verify whether (Up, Right) is played but cannot distinguish among any of the other action profiles. Thus, a different transfer τ can be enforced for (Up, Right), but all other action profiles share the same transfer τ'.

Suppose for the moment that $\delta = 0$. The profile (Down, Left) gives the highest joint value but is not a Nash equilibrium for any τ and τ'. For $\tau - \tau' > -1$ there is a mixed-strategy equilibrium in which the players choose Out with probability 0, player 1 chooses Up with probability $1/(2 + \tau - \tau')$, and player 2 chooses Left with probability $(1 + \tau - \tau')/(3 + \tau - \tau')$. In this equilibrium, (Down, Left) is played with a probability that is increasing in $\tau - \tau'$. There is no maximum equilibrium joint value by choice of $\tau, \tau' \in \mathbb{R}$ and therefore we cannot guarantee existence without restricting the class of stage games, such as bounding transfers. The problem extends to the setting with $\delta > 0$.

Overall, bounding transfers may help secure equilibrium existence but, for stage games like the one above, bounds interfere with our main result. This is due to a trade-off between using constant transfers to provide incentives in previous periods and using differential transfers to provide incentives in the current period. For example, suppose τ and τ' must be in $[-4, 4]$ and assume δ is strictly positive but small enough so that cooperation still requires a mixed action profile. The players, in agreement, want an external contract that specifies $\tau = 4$ and $\tau' = -4$ in the current period, for this gives the best incentives for the stage game. They would also like to pick a continuation external contract from the next period with a continuation value that favors player 1 in the event of (Up, Right) and favors player 2 otherwise. But to do this, the players would want τ' in the next period to be larger than -4 following (Up, Right) in the current period. This would generally not maximize the span from the next period, however, and the current-period transfer constraints do not allow an adjustment to utilize the maximal-span continuation contract.

B3. *Noncooperative Foundations*

Our hybrid cooperative/noncooperative model is tightly connected to a fully noncooperative account of the contractual setting in which the negotiation phase is described as a bargaining protocol, such as random-proposer ultimatum-offer. Watson (2013) and Miller and Watson (2013) develop a refinement of perfect public equilibrium based on axioms that relate statements and voluntary transfers in the bargaining phase to a selection of continuation play from the action phase in the current period.[37] The refinement, called contractual equilibrium in the fully noncooperative game, is equivalent to the recursive formulation of contractual-equilibrium continuation values.

The Miller and Watson (2013) analysis extends with minimal modification to our setting with external enforcement. An offer includes (i) a contract c, (ii) an immediate transfer, and (iii) a specification of future behavior summarized by continuation values. Acceptance of an offer causes c to be externally enforced and causes the immediate transfer to be automatically enforced as well (not necessarily by the same authority that enforces c). Axioms relate the third part of the offer to the coordinated play in the continuation of the game.[38] External enforcement adds one new technicality, related to the existence issue described in Section B1: it is feasible for the players to enter the action phase of a period with a contract c (by default or by agreement) for which there is no equilibrium action profile. To deal with this problem in general, one can ignore the equilibrium conditions for such contingencies or limit Γ to finite stage games (where the problem would not arise). A failure of joint-value maximization, as in the second example described in B2, would lead to nonexistence, just as in the hybrid model.

REFERENCES

Abreu, Dilip, David G. Pearce, and Ennio Stacchetti. 1990. "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring." *Econometrica* 58 (5): 1041–63.

Aghion, Philippe, Mathias Dewatripont, and Patrick Rey. 1994. "Renegotiation Design with Unverifiable Information." *Econometrica* 62 (2): 257–82.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 1994. "Subjective Performance Measures in Optimal Incentive Contracts." *Quarterly Journal of Economics* 109 (4): 1125–56.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 2002. "Relational Contracts and the Theory of the Firm." *Quarterly Journal of Economics* 117 (1): 39–84.

Baker, George, Robert Gibbons, and Kevin J. Murphy. 2011. "Relational Adaptation." Unpublished.

Barron, Daniel, Robert Gibbons, Ricard Gil, and Kevin J. Murphy. 2019. "Relational Adaptation under Reel Authority." *Management Science* 66 (5): 1868–89.

Bernheim, B. Douglas, and Michael D. Whinston. 1998. "Incomplete Contracts and Strategic Ambiguity." *American Economic Review* 88 (4): 902–32.

Beuve, Jean, and Stéphane Saussier. 2012. "Interfirm Cooperation in Strategic Relationships: The Role of Formal Contract." *Industrial and Corporate Change* 21 (4): 811–36.

Blair, Roger D., and Francine Lafontaine. 2005. *The Economics of Franchising.* Cambridge: Cambridge University Press.

Bolton, Patrick, and Mathias Dewatripont. 2005. *Contract Theory.* Cambridge, MA: MIT Press.

Bull, Clive. 1987. "The Existence of Self-Enforcing Implicit Contracts." *Quarterly Journal of Economics* 102 (1): 147–59.

Buzard, Kristy, and Joel Watson. 2012. "Contract, Renegotiation, and Holdup: Results on the Technology of Trade and Investment." *Theoretical Economics* 7: 283–322.

[37] In Miller and Watson (2013), an internal agreement axiom requires that if the players agree to a continuation that is incentive-compatible and consistent with their current equilibrium from the next period, then they will play as agreed. A no-fault disagreement condition requires that, in a disagreement outcome of the bargaining process, continuation play does not depend on how disagreement occurred, and failing to make a promised immediate transfer constitutes disagreement. Finally, an external agreement condition requires that in negotiation, the players jointly optimize over equilibria that satisfy the first two axioms.

[38] The agreement axiom applies to agreements about future play that are feasible given the selected c. Note that we assume immediate transfers are automatically enforced, so they are tied to the selection of c, as in Watson (2013). This is important for the equilibrium characterization. Without external enforcement, transfers can be voluntary and occur after players voice agreement, as in Miller and Watson (2013).

Che, Yeon-Koo, and Donald B. Hausch. 1999. "Cooperative Investments and the Value of Contracting." *American Economic Review* 89 (1): 125–47.

Che, Yeon-Koo, and Seung-Weon Yoo. 2001. "Optimal Incentives for Teams." *American Economic Review* 91 (3): 525–41.

Dalen, Dag Morten, Espen R. Moen, and Christian Riis. 2006. "Contract Renewal and Incentives in Public Procurement." *International Journal of Industrial Organization* 24 (2): 269–85.

Edlin, Aaron S., and Stefan Reichelstein. 1996. "Holdups, Standard Breach Remedies, and Optimal Investment." *American Economic Review* 86 (3): 478–501.

Evans, Robert. 2008. "Simple Efficient Contracts in Complex Environments." *Econometrica* 76 (3): 459–91.

Fong, Yuk-fai, and Jin Li. 2017. "Relational Contracts, Limited Liability, and Employment Dynamics." *Journal of Economic Theory* 169: 270–93.

Goldlücke, Susanne, and Sebastian Kranz. 2013. "Renegotiation-Proof Relational Contracts." *Games and Economic Behavior* 80: 157–78.

Hart, Oliver, and John Moore. 1988. "Incomplete Contracts and Renegotiation." *Econometrica* 56 (4): 755–85.

Hart, Oliver, and John Moore. 1999. "Foundations of Incomplete Contracts." *Review of Economic Studies* 66 (1): 115–38.

Iossa, Elisabetta, and Giancarlo Spagnolo. 2011. "Contracts as Threats: On a Rationale for Rewarding A while Hoping for B." Unpublished.

Itoh, Hideshi, and Hodaka Morita. 2015. "Formal Contracts, Relational Contracts, and the Threat-Point Effect." *American Economic Journal: Microeconomics* 7 (3): 318–46.

Johnson, Simon, John McMillan, and Christopher Woodruff. 2002. "Courts and Relational Contracts." *Journal of Law, Economics, and Organization* 18 (1): 221–77.

Klein, Benjamin, and Keith B. Leffler. 1981. "The Role of Market Forces in Assuring Contractual Performance." *Journal of Political Economy* 89 (4): 615–41.

Kostadinov, Rumen. 2019. "Renegotiation of Long-Term Contracts as Part of an Implicit Agreement." Unpublished.

Kvaløy, Ola, and Trond E. Olsen. 2009. "Endogenous Verifiability and Relational Contracting." *American Economic Review* 99 (5): 2193–208.

Lazzarini, Sergio G., Gary J. Miller, and Todd R. Zenger. 2004. "Order with Some Law: Complementarity versus Substitution of Formal and Informal Arrangements." *Journal of Law, Economics, and Organization* 20 (2): 261–98.

Levin, Jonathan. 2003. "Relational Incentive Contracts." *American Economic Review* 93 (3): 835–57.

Macaulay, Stewart. 1963. "Non-Contractual Relations in Business: A Preliminary Study." *American Sociological Review* 28 (1): 55–67.

MacLeod, W. Bentley, and James M. Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica* 57 (2): 447–80.

Macneil, Ian R. 1978. "Contracts: Adjustment of Long-Term Economic Relations under Classical, Neoclassical, and Relational Contract Law." *Northwestern University Law Review* 72 (6): 854–904.

Malcomson, James M. 2013. "Relational Incentive Contracts." In *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, 1014–65. Princeton, NJ: Princeton University Press.

Martimort, David, Aggey Semenov, and Lars Stole. 2017. "A Theory of Contracts with Limited Enforcement." *Review of Economic Studies* 84 (2): 816–52.

Maskin, Eric, and Jean Tirole. 1999. "Unforeseen Contingencies and Incomplete Contracts." *Review of Economic Studies* 66 (1): 83–114.

Miller, David A., and Joel Watson. 2013. "A Theory of Disagreement in Repeated Games with Bargaining." *Econometrica* 81 (6): 2303–50.

Narasimhan, Subha. 1989. "Relationship or Boundary? Handling Successive Contracts." *California Law Review* 77 (5): 1077–122.

Nash, John F., Jr. 1950. "The Bargaining Problem." *Econometrica* 18 (2): 155–62.

Nöldeke, Georg, and Klaus M. Schmidt. 1995. "Option Contracts and Renegotiation: A Solution to the Hold-Up Problem." *Rand Journal of Economics* 26 (2): 163–79.

Poppo, Laura, and Todd Zenger. 2002. "Do Formal Contracts and Relational Governance Function as Substitutes or Complements?" *Strategic Management Journal* 23 (8): 707–25.

Ray, Debraj. 2002. "The Time Structure of Self-Enforcing Agreements." *Econometrica* 70 (2): 547–82.

Ryall, Michael D., and Rachelle C. Sampson. 2009. "Formal Contracts in the Presence of Relational Enforcement Mechanisms: Evidence from Technology Development Projects." *Management Science* 55 (6): 906–25.

Safronov, Mikhail, and Bruno Strulovici. 2018. "Contestable Norms." Unpublished.
Schmidt, Klaus M., and Monika Schnitzer. 1995. "The Interaction of Explicit and Implicit Contracts." *Economics Letters* 48 (2): 193–99.
Segal, Ilya. 1999. "Complexity and Renegotiation: A Foundation for Incomplete Contracts." *Review of Economic Studies* 66 (1): 57–82.
Shapiro, Carl, and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74 (3): 433–44.
Watson, Joel. 2001. *Strategy: An Introduction to Game Theory.* 1st ed. New York: WW Norton & Co.
Watson, Joel. 2007. "Contract, Mechanism Design, and Technological Detail." *Econometrica* 75 (1): 55–81.
Watson, Joel. 2013. "Contract and Game Theory: Basic Concepts for Settings with Finite Horizons." *Games* 4 (3): 457–96.