

UC San Diego

UC San Diego Previously Published Works

Title

Theoretical Foundations of Relational Incentive Contracts

Permalink

<https://escholarship.org/uc/item/19f9w2xf>

Journal

Annual Review of Economics, 13(1)

ISSN

1941-1383

Author

Watson, Joel

Publication Date

2021-08-05

DOI

10.1146/annurev-economics-090820-110736

Peer reviewed

Theoretical Foundations of Relational Incentive Contracts*

Joel Watson[†]

January 2021

Abstract: This article describes the emerging game-theoretic framework for modeling long-term contractual relationships with moral hazard. The framework combines self-enforcement and external enforcement, accommodating alternative assumptions regarding how actively the parties initially set and renegotiate the terms of their contract. A progression of theoretical components is reviewed, building from the recursive formulation of equilibrium continuation values in repeated games. A principal-agent setting serves as a running example. *Keywords: relational contracts, enforcement, game theory, negotiation.*

A *contract* is an agreement that is intended to be enforced. Almost all contracts are enforced with a combination of *external enforcement* and *self-enforcement*. External enforcement refers to actions of third parties, such as courts and other legal authorities collectively called the *external enforcer*, that influence the contracting parties' behavior. Self-enforcement entails coordinated actions that the contracting parties themselves take, consistent with their individual incentives. In a *relational-contracting environment*, the parties interact over time, so there is an intertemporal aspect of self-enforcement whereby behavior at one time relates to the parties' anticipated response later.¹

Relational contracts are everywhere and in every branch of economics. The scientific literature on the topic is dispersed. A subset of the relational-contract literature sometimes called "relational incentive contracts" has congregated on a class of repeated-game style models of ongoing relationships with moral hazard, most notably ones involving a firm's arrangements with suppliers or its management of workers. Models in this area vary in

*This is a short version of an overview of relational contracting in the author's forthcoming manuscript on contract theory.

[†]<https://econweb.ucsd.edu/~jwatson/>. The author thanks Trond Olsen, Xiameng Hua, Stephen Morris, and Matthias Fahn for their generous input.

¹An organized society contracts in multitude layers, and at some level all enforcement is self-enforcement. But in models of individual relationships, it is convenient to treat third parties as external and to abstract from their incentives.

terms of time lines, technological assumptions, and formalism. As a result, it may be challenging for prospective researchers to find the modeling elements essential for capturing the literature's broad themes and insights.

This article gives an overview of a core game-theoretic framework that has emerged in recent years for the analysis of relational incentive contracts. The framework covers relationships with the following attributes:

- discrete-time interaction with an infinite number of periods;
- external enforcement of monetary transfers contingent on verifiable information;
- a stationary productive environment, without initial private information;
- multiple-phase interaction within a period, with monetary transfers and negotiation in one phase and productive interaction in another; and
- payoffs that are linear in money.

Equilibrium analysis is put in terms of dynamic programming, where equilibrium continuation values are characterized recursively. Ultimate versions of the framework incorporate bargaining theory to model active contracting.

To describe the core theory, this article reviews in linear fashion a small number of key technical steps taken in the literature. I focus on fundamental definitions and analytical methods, without technical details and formal results. The concepts and methods are illustrated with a running numerical example. I assume the reader has knowledge of game theory at the level of a solid game-theory textbook such as Watson (2002).

I hope this presentation will be useful to readers in multiple categories. For the applied researcher, it may serve as a bridge to the theory and essential techniques. For graduate students and other scholars, it introduces the theory of relational contracts and details the special structure of these game-theoretic models. The modeling framework may be a useful starting point for future theoretical studies.

Terminology

In preparation for the technical presentation, here are some notes on the terminology used in the contract-theory literature. Because the models are game-theoretic, we typically call the contracting parties *players* and their utilities *payoffs*. In addition to the distinction between self-enforcement and external enforcement, there are important distinctions to be made regarding information and the components of contract. *Private information* refers to something that one player observes but other players do not observe. Items that all of the players see are called *commonly observed*, and information available to an external enforcer is called *verifiable*.²

Every contract can be described as combining an external part and an internal part. *External contract* refers to the contractual provisions that instruct the external enforcer on how to intervene in the relationship, typically by compelling monetary transfers as a function of verifiable information. The internal part of a contract records how the contracting

²Foundations of verifiability are developed by Bull and Watson (2004).

parties have agreed to themselves act. Note that this terminology differs from other terms used in various strands of the literature. For instance, lawyers traditionally reserved the term “contract” for an agreement, ideally formalized in writing, that the court would enforce. As a result, when distinguishing between external and self-enforcement, scholars have adopted terms like “formal versus informal” and “explicit versus implicit.” But these words create confusion, for plenty of agreements that could not be enforced by courts are nonetheless highly detailed and formalized in writing. Likewise, many contracts that rely on court enforcement are informally expressed. Therefore, I stick to the external/internal terminology.

Historical notes

The term “relational contract” originated in the legal literature, where scholars observed the prevalence of self-enforcement in long-term relationships, studied examples, and discussed the relation to external enforcement. Macaulay (1963) and Macneil (1978) are focal references. The starting point in the formal analysis of relational contracts by economists and game theorists is not clear. But as the mechanics of self-enforcement in relational contracts are given by the conditions for equilibria in repeated games, the core elements of relational contracting were first developed in the late 1950s by game theorists. Rubinstein (1979) examined a repeated-game model of a principal-agent relationship with binary choices for the two parties, and the following years saw the introduction of applied models with monetary transfers. Telser (1980) and Klein and Leffler (1981) modeled repeat purchases, where prices above the competitive level give firms the incentive to provide high quality over time. Shapiro and Stiglitz (1984) examined employment relationships, similarly finding that high “efficiency wages” induce workers to exert effort. Bull (1987) distinguished between self-enforcement and external enforcement in a finite-period model.

In the large literature that developed since the early 1980s, a few articles stand out as exemplars with regard to technical steps that, collectively, led to the general framework that I present here. Radner (1985) developed a discrete-time, repeated-game model of a principal-agent setting with an infinite horizon and external enforcement of short-term contracts (monetary transfers contingent on output, assumed observable and verifiable). He studied a class of equilibria in which the players use “review strategies.” Spear and Srivastava (1987) examined equilibria more generally and put the analysis in terms of a dynamic program featuring continuation values. Meanwhile, MacLeod and Malcolmson (1989) developed a model of employment relationships that distinguishes between observable and verifiable aspects of production (they assumed that only employment is verifiable) and with separate phases within a period for transfers and the worker’s effort choice. This led to a more general analysis of contractual arrangements than Shapiro and Stiglitz (1984) studied. The assumption of quasilinear utility, in both papers, simplified the analysis of incentive conditions.

The full-blown recursive characterization of “perfect public equilibrium” values in general repeated games was pioneered by Abreu, Pearce, and Stacchetti (1990). Later Goldlücke and Kranz (2012, 2013) characterized perfect public equilibrium values for gen-

eral repeated games with quasilinear utility and separate phases for transfers and productive actions in each period. Miller and Watson (2013) added an explicit account of bargaining and defined “contractual equilibrium” in general settings with self-enforcement. Finally Watson, Miller, and Olsen (2020) extended the theory to settings with external enforcement of long-term contracts.³

Popular theoretical applications

To appreciate the generality of the core theory presented here, it is useful to note the range of theoretical applications that appear in the recent literature. Applications vary in terms of the details of production and the presence of additional strategic elements. I mention a few popular ones next, citing one or two representative papers for each (by no means a comprehensive list). Additional citations appear later in this article. For a broader discussion of applications and economic insights, the reader may want to look also at MacLeod (2007) and Malcomson (2013).

The most extensively studied settings are long-term principal-agent relationships, particularly employment relations (as in MacLeod and Malcomson 1989 and Levin 2001) but also a variety of similar relationships such as between regulators and firms (Bertelli and Smith 2010) or environmental organizations and communities (Gjertsen et. al. 2020). Multi-sided moral hazard is present in models of production by teams within a firm, partnerships, and joint-ventures (Doornik 2006) and also international agreements (Klimenko, Ramey, and Watson 2008). Models of supply management study managers’ policies to select from qualified suppliers over time (Board 2011, Andrews and Barron 2016).

A great deal of research has been aimed at explaining specific employment practices and the form of contracts, in some cases extending the core theory to include elements such as private information. Influential early entries include Holmstrom and Milgrom (1987) on a justification for linear bonus contracts; Baker, Gibbons, and Murphy (1999) on the source of authority in organizations; Baker, Gibbons, and Murphy (1994, 2002) and Schmidt and Schnitzer (1995) on the interaction between external and internal contracts; and Halonen (2002) on ownership structure in relational contracts. More recent entries include analysis of how routines emerge in organizations (Chassang 2010), dynamics related to restricted transfers (Li, Matouschek, and Powell 2017), and private monitoring with verifiable reports (Fuchs 2007). More distant from the core theory presented here are studies of long-term asymmetric information and project choice (Watson 1999, 2002) and privately known outside options (Halac 2012).

³Here are other notable entries that I will discuss later: Levin (2001) expanded MacLeod and Malcomson’s (1989) model to study a range of production technologies and to further characterize equilibrium strategies. Baker, Gibbons, and Murphy (1994, 2002) and Schmidt and Schnitzer (1995) explored the interaction of external and self-enforcement in settings where parties have more information than is verifiable. Ramey and Watson (1997, 2001) and den Haan, Ramey, and Watson (2000) put relational contracts in the context of a matching market, showing how incentives in employment relationship interact across markets in the presence of shocks.

Outline

I begin in the next section with the description of a short-term contractual setting, to show the conceptual connections between short-term and long-term relationships. Section 2 describes contractual settings with an infinite horizon, where the phases of the short-term setting repeat but where the external contract may create nonstationarities. This environment has all of the key features highlighted in the brief historical review above. Section 2 ends with a description of the operator used for the recursive characterization of equilibrium values. Section 3 provides details of the framework for settings with trivial external enforcement, where the focus is on self-enforcement, and Section 4 expands the framework to settings with nontrivial external enforcement.

1 Setting the Stage: Modeling a Short-Term Relationship

Let us begin with a contractual setting in which all productive actions take place at one time (a static setting). There are n contracting parties, also called *players*. They interact first in the *negotiation phase*, where they negotiate a contract and make immediate monetary transfers, and then in the *production phase* (also called the *productive-action phase*), where they simultaneously choose productive actions and receive payoffs. The production phase is formally described as a *stage game*.

The players' contract combines an *external part* and an *internal part*. The external part specifies how the external enforcement authority should intervene in the relationship, such as by compelling monetary transfers as a function of verifiable information or providing a monitoring service. All externally enforced aspects of the contract are assumed to be represented in the stage game, which I denote by γ . Thus, in the negotiation phase, the players essentially pick a stage game γ from a set of feasible games Γ , where Γ represents the scope of external enforcement.⁴ The internal part of the contract specifies how the players have agreed to behave in the stage game; this part must be self-enforced.⁵

Behavior in the production phase is modelled noncooperatively, where self-enforcement is typically expressed as a Nash equilibrium of the stage game. Behavior in the negotiation phase can also be modelled noncooperatively, with a bargaining protocol, or by using a cooperative solution concept such as the generalized Nash bargaining solution. Let us take the latter route and denote by $\pi = (\pi_1, \pi_2, \dots, \pi_n)$ the fixed vector of nonnegative bargaining weights, where $\pi_1 + \pi_2 + \dots + \pi_n = 1$.⁶

⁴The parties' choice of a production technology can also be incorporated into the set Γ .

⁵How the theorist's abstract model relates to contracting and enforcement in the real world is an important topic, especially because actual enforcement authorities are not as passive as our theoretical models typically assume. For an overview in the context of the basic setting described below, see Watson (2002, Chapter 13). For discussion and analysis of contract interpretation, see Shavell (2006), Listokin (2010), and Schwartz and Watson (2013).

⁶If we model the negotiation phase noncooperatively, then we are looking at a grand noncooperative game consisting of a bargaining protocol in the negotiation phase followed by the stage game in the second phase. Under some assumptions regarding how the players coordinate their behavior in the stage game, parameters

1.1 The basic setting

Let us focus on a special case along the lines of Holmström (1982) and Legros and Matthews (1993): utility that is linear in money, a fixed production technology, and external enforcement of only monetary transfers. In this setting, the technology of productive interaction is fixed and described by an *underlying game*, representing the players' productive actions, personal costs and benefits, and the intrinsic distribution of returns.⁷ The underlying game is denoted $\gamma = (A, X, \lambda, u, P)$, with the components described as follows:

- a set of action profiles $A = A_1 \times A_2 \times \cdots \times A_n$,
- an outcome set X ,
- a conditional distribution function $\lambda: A \rightarrow \Delta X$,
- a payoff function $u: A \rightarrow \mathbb{R}^n$, and
- a partition P of X representing verifiability constraints.

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.⁸ 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 , 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 .⁹

The external part of the players' contract specifies an externally enforced monetary transfer between them as a function of the outcome x . It is a function $b: X \rightarrow \mathbb{R}_0^n$, where

$$\mathbb{R}_0^n = \left\{ m \in \mathbb{R}^n \mid \sum_{i=1}^n m_i = 0 \right\}$$

is the space of balanced transfers. For any action profile a and outcome x in the underlying game, the payoff vector is $u(a) + b(x)$. Let $\bar{b}(a) \equiv E_{\lambda(a)} b(x)$ be the expected transfer given action profile $a \in A$. Then the externally enforced transfer transforms the underlying game into the *induced game* given by

$$\langle A, u + \bar{b} \rangle, \tag{1}$$

and this is the stage game that the players effectively play in the production phase.

Importantly, transfer function b is constrained to be P -measurable, because the external enforcer can observe only what is verifiable about the outcome. Thus, if the external enforcer cannot distinguish between outcomes x and x' , meaning that $x \in P(x')$, then

of the bargaining protocol translate into the bargaining weights in the hybrid model where the negotiation phase is modelled cooperatively.

⁷Ruled out here are elements such as decisions about production technology and third-party monitoring services.

⁸" ΔX " denotes the set of probability distributions over X .

⁹To 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. This framework also accommodates applications in which the players may not observe their own payoffs.

$b(x) = b(x')$ is required. Note also that, by assuming only balanced transfers are enforced, the model does not allow the players to commit to send money to a third party or otherwise throw away money. The justification for such an assumption is that if the players had specified such a thing, then once the outcome occurs they would have the joint incentive to renegotiate their contract in order to release themselves from this obligation.

In summary, the contractual setting plays out as follows: In the negotiation phase, the players form a contract specifying an immediate transfer, a transfer function b , and a mixed action profile $\alpha \in \Delta A$ (an uncorrelated probability distributions over A) for the production phase. This leads to the induced game in Expression 1, where the players have agreed to coordinate on action profile α .¹⁰ Self-enforcement requires α to be a Nash equilibrium of the induced game, so that each player best-responds to the others' action profile.

To round out the model, we must describe the disagreement point of bargaining, which is what would happen if the players fail to make an agreement in the negotiation phase. Suppose that disagreement involves no immediate transfer and $b \equiv \mathbf{0}$, the constant function that gives a transfer of zero, so that the stage game is simply the unaltered underlying game. Also suppose that the players would then coordinate on an exogenously given Nash equilibrium $\underline{\alpha}$ of the underlying game.¹¹

With transferable utility and efficient bargaining, as the Nash solution predicts, clearly in the negotiation phase the players will choose the external and internal contractual elements b and α to maximize their joint value $\sum_{i=1}^n u_i(a)$ subject to b being P -measurable and α being a Nash equilibrium of the induced game in Expression 1. Letting α^* and b^* denote a solution, and defining $L^* \equiv \sum_{i=1}^n u_i(\alpha^*)$ and $\underline{w} \equiv u(\underline{\alpha})$, the Nash bargaining solution predicts that the players will make an up-front transfer to achieve the payoff vector $\underline{w} + \pi (L^* - \sum_{i=1}^n \underline{w}_i)$ from the negotiation phase.

It is easy to see that improvements in the external enforcement technology, such as increased verifiability as represented by a more refined partition P , can only improve the prospects for aligning incentives and must weakly increase welfare.

1.2 Example

I next introduce a two-player example that will be used throughout this article. Consider a relationship between a worker (player 1) and a manager (player 2). In the underlying game, the worker chooses whether to exert effort and, if so, to which of three projects to apply his effort. He can expend effort on only one project. The manager observes the worker's effort choice and receives the revenue that it generates. The manager has no action in the underlying game. The worker's effort choice is unverifiable, but the outcome includes a

¹⁰It is not necessary at this point to separate the immediate transfer from the outcome-contingent transfer. The former can be incorporated in b . However, the added terms helps organize our accounting of payoffs and continuation values in multi-period settings.

¹¹Subtle issues arise here that are sometimes swept under the rug in the contract-theory literature. For instance, what would we predict if the underlying game has no Nash equilibrium or multiple Nash equilibria? Further, what would happen in other off-equilibrium-path contingencies, such as if the players made an agreement other than the one our theory predicts?

a_1	player 1's cost	player 2's revenue	$\sigma(a_1)$
0	0	0	0
1	11	19	1/2
2	1	7	1/4
3	22	28	1

Table 1: Payoff and signal parameters in the project-choice example.

verifiable binary signal of this action.

The set of feasible effort choices is $A_1 = \{0, 1, 2, 3\}$, where $a_1 = 0$ represents no effort and $a_1 > 0$ means applying effort to project a_1 . The signal is 1 with probability $\sigma(a_1)$ and 0 with probability $1 - \sigma(a_1)$. For each effort level, the Table 1 gives the worker's effort cost, the manager's revenue, and the probability of the high signal. Note that $a_1 = 1$ is the efficient effort choice, yielding a joint value of 8. Effort choices 2 and 3 each yields a joint value of 6. The choices all differ in terms of the probability of the high signal. The payoff vectors for the underlying game, along with the frontier of feasible values utilizing transfers, are pictured in the left graph of Figure 1.

The outcome space is $X \equiv \{00, 01, 10, 11, 20, 21, 30, 31\}$, where the first digit of the outcome is a_1 and the second digit is the realization of the signal. The contingent distribution function λ is given by $\lambda(0)(00) = 1$, $\lambda(1)(10) = 1/2$, $\lambda(1)(11) = 1/2$, $\lambda(2)(20) = 3/4$, $\lambda(2)(21) = 1/4$, and $\lambda(3)(31) = 1$. Assume that the signal is verifiable but player 1's effort choice is not verifiable, so the outcome partition is

$$P = \{\{00, 10, 20, 30\}, \{01, 11, 21, 31\}\}.$$

In this setting, the external contract b essentially specifies a bonus ρ to be transferred from player 2 to player 1 in the event of the high signal, along with a constant baseline transfer that we can set to zero without loss of generality. The middle graph of Figure 1 shows the induced game for a contract specifying $\rho = 4$, while the right graph shows the induced




Figure 1: Payoffs in the project-choice example.

game corresponding to $\rho = 28$. Note the change in scale relative to the left graph.

Because only player 1 has an action in the stage game, it is easy to visualize player 1's incentives by looking at the graph of the payoffs in the induced game. Clearly, player 1 will select whatever action corresponds to the right-most point in the graph, yielding the highest payoff for this player. It is thus helpful to consider the "implementation problem" where, for any given a_1 , we determine whether there is a value of ρ that would give player 1 the incentive to choose this action—that is, that makes $u_1(a_1) + \bar{b}(a_1)$ the right-most point in the graph.

There are four things to note for this example. First, the players can achieve a joint value of 6 by agreeing to a bonus that satisfies $\rho \geq 4$, such as those illustrated in the middle and right graphs of Figure 1. Second, the players cannot do better because no contract can implement the efficient action $a_1 = 1$. This is easy to see by observing that $u_1(2) + \bar{b}(2) > u_1(1) + \bar{b}(1)$ for $\rho < 40$, and $u_1(3) + \bar{b}(3) > u_1(1) + \bar{b}(1)$ for $\rho > 22$.

Third, the difference

$$\max_{a_1=0,2,3} u_1(a_1) + \bar{b}(a_1) - (u_1(1) + \bar{b}(1)),$$

though strictly positive, is minimized by choosing $\rho = 28$ as shown in the right graph of Figure 1. In other words, $\rho = 28$ provides the greatest incentive for player 1 to choose $a_1 = 1$, but it is still not enough to motivate player 1 to actually choose this action. An implication is that, if we were to embellish the model with a complementary element to enhance the incentive for player 1 to choose $a_1 = 1$, then it would be best for the contract specify $\rho = 28$. Fourth, note that if the players' contract specifies $\rho = 4$, then the induced game has two Nash equilibria, $a_1 = 0$ and $a_1 = 2$. The third and fourth items will play an important role in the analysis of relational contracting that I turn to next.

2 Repeated Interaction and Relational Contracts

In a relational-contract setting, productive interaction occurs over multiple periods of time. In this section I describe the main components of a general model of relational contracting with an infinite number of discrete periods. I begin by describing the payoff relevant elements and discuss how one can organize equilibrium analysis using the notion of *continuation value*, which is the payoff vector in the continuation of the game from any particular point in time. I then describe the key operator that relates sets of equilibrium continuation values across periods.

2.1 Main ingredients of the basic model

The basic relational-contracting model has an infinite time horizon, discrete periods, and a fixed production technology that the players engage in repeatedly. The time period is denoted $t = 1, 2, 3, \dots$. In each period, the players interact just as in the short-term setting: the negotiation phase followed by the production phase. In the production phase, the

players choose actions in an underlying game (A, X, λ, u, P) , which I now sometimes call the *underlying stage game*, and the outcome determines a monetary transfer according to the function $b: X \rightarrow \mathbb{R}_0^n$ that their external contract specifies. These elements are defined exactly as before. The underlying stage game is fixed and the same in all periods.

To incorporate discounting, assume that the players have a common discount factor $\delta \in (0, 1)$ that, by multiplication, translates a payoff received in any period $t + 1$ into the period- t equivalent. Further, let us normalize payoffs in each period by multiplying by $1 - \delta$, which will be useful later to facilitate comparisons between stage-game payoffs and continuation values. This is a standard normalization in repeated-game theory.

The external contract is in general more complicated than in the single-period setting. It specifies a transfer function $b^t: X \rightarrow \mathbb{R}_0^n$ for each period t , which gives the monetary transfer at the end of this period as a function of this period's outcome. The contract could specify different transfer functions for different periods. Further, it could make b^t a function of the verifiable history of the relationship through period $t - 1$. For instance, in the project-choice example, the contract can increase or decrease the bonus in period t based on the values of the signal in previous periods.

The model may be starting to sound complicated, but bear with me because I'll next scale back the generality to describe the simple setting in which the players are restricted to contractual provisions that make b^t the same in each period (equivalently, where the external enforcement authority will enforce only contracts of this form). Focusing on this restricted "stationary" setting will allow me to provide intuition and to describe the recursive methodology for analyzing relational contracting problems. Further, you'll be happy to know that a large part of the relational contracting literature focuses on such restricted settings. But don't get too elated, because this restriction is unrealistically artificial and I'll return to the more general setting later in this article, for currently the most active areas of relational-contract analysis feature elements that make $\{b^t\}$ nonstationary.

I should also note at this point that some standard relational-contracting models incorporate additional strategic elements within a period of time. One such element is an outside-option phase at the end of the period, where the players simultaneously decide whether to continue or end their relationship. If one or both players elects to sever the relationship, then they receive terminal payoffs that represent their values of finding other trading partners or working on their own. Models sometimes also include a voluntary-transfer phase between the production phase and the outside-option phase. I will comment on these variations later.

As with the short-term setting, there are multiple ways to model behavior in the negotiation phase of each period. There is the choice of cooperative or noncooperative bargaining models, and there is a variety of alternative assumptions one can make regarding how communication in the negotiation phase translates into coordinated play in the production phase of the current period and beyond. I will leave the negotiation phase unspecified for now, but note that the payoff relevant aspect of play is an immediate transfer the players make in reaching an agreement. Let us denote by m^t the transfer in the negotiation phase of

period t . I assume $m^t \in \mathbb{R}_-^n$ where

$$\mathbb{R}_-^n \equiv \{m \in \mathbb{R}^n \mid m_1 + m_2 + \dots + m^n \leq 0\}.$$

That is, the players can transfer money between them and also can throw away money.

2.2 Continuation values and equilibrium

There are two standard approaches to characterizing equilibria in repeated games and relational-contract models. The first involves describing strategies for the entire game and then stating and evaluating equilibrium conditions on the strategy space. The second approach, pioneered by Abreu, Pearce, and Stacchetti (1990), characterizes the set of equilibrium continuation values recursively via dynamic programming. I focus on the recursive approach here, because it is particularly elegant and has a special structure due to special features of relational contracting models. This subsection initiates the analysis with a note on equilibrium concepts and the definition of continuation values.

The solution concepts we will use assume that equilibrium behavior in each period depends on only the jointly observable history of play. In particular, in equilibrium a player does not condition her behavior on her own past productive actions except to the extent that they are revealed by the outcome. This is the standard restriction that defines the notion of *perfect public equilibrium* in repeated-game theory.

The payoff vector in period τ is the expectation of $(1 - \delta)(m^\tau + u(a^\tau) + b^\tau(x^\tau))$, where m^τ is the transfer made in the negotiation phase, a^τ is the action profile played in the production phase, x^τ is the outcome, and b^τ is the transfer function in period τ . The continuation value from the beginning of any period t is the expected discounted sum of the payoff vectors from this period on. Denoting this value y^t , we have:

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

where the expectation is taken with respect to the distribution of $\{m^\tau, a^\tau, x^\tau, b^\tau\}_{\tau=t}^{\infty}$.

The continuation values we need to analyze are those that arise in equilibrium. Ignoring that I have not finished describing the game or the equilibrium concept, suppose that for a given history of play through period $t - 1$, the equilibrium will specify behavior in the continuation of the game from period t , resulting in a continuation value y^t . We can write y^t as the sum of the expected payoff in period t and the discounted continuation value from period $t + 1$:

$$y^t = E [(1 - \delta) (m^t + u(a^t) + b^t(x^t)) + \delta y^{t+1}],$$

where the expectation is conditioned on the history of play through period $t - 1$.

We can describe the continuation value from the production phase similarly. Fix the history of play through period $t - 1$, and also fix m^t and b^t . Define $\bar{b}^t : A \rightarrow \mathbb{R}_0^n$ to give the expectation of b^t as a function of the action profile, as in the short-term setting. Suppose that the players are about to interact in the production phase of period t . If action profile

$a^t \in A$ is played then the players obtain $(1 - \delta)(u(a^t) + \bar{b}^t(a^t))$ in the current period and continuation value y^{t+1} from the start of the next period. Since all of this is already conditioned on the history through the negotiation phase of period t , we can think of y^{t+1} as a function only of x^t , the verifiable outcome of production in period t . We can then write the continuation value from the production phase as a function of the action profile a^t chosen in the current period. To do this, define $\bar{y}^{t+1}: A \rightarrow \mathbb{R}^2$ by $\bar{y}^{t+1}(a^t) \equiv E_{\lambda(a^t)} y^{t+1}(x^t)$ for every $a^t \in A$. Therefore, if a^t is played in period t , the continuation value from the production phase is

$$(1 - \delta) \left(u(a^t) + \bar{b}^t(a^t) \right) + \delta \bar{y}^{t+1}(a^t). \quad (2)$$

2.3 Incentives in the production phase of a single period

We next explore the players' incentives in the production phase of any given period t . Rather than refer to period t by number, I'll now refer to it as "the current period." Dropping the superscripts in Expression 2, let b be the contracted transfer function for the current period, let y denote the continuation value from the start of the next period, and let \bar{b} and \bar{y} give the expectations of b and y as a function of the action profile. Then Expression 2 becomes $(1 - \delta)(u(a) + \bar{b}(a)) + \delta \bar{y}(a)$, where a is the action profile in the current period.

Therefore, interaction in the production phase of the current period is essentially play of the induced game

$$\langle A, (1 - \delta)(u + \bar{b}) + \delta \bar{y} \rangle. \quad (3)$$

Self-enforcement in the current period amounts to coordination on a Nash equilibrium of this induced game. In comparison to the induced game in Expression 1 in the short-term setting, the induced game here simply adds the continuation value as an extra consequence of the productive actions. As in the short-term setting, for various functions b and y , we can determine whether any given mixed action profile α is self-enforced in the current period as a Nash equilibrium of induced game, and calculate the resulting payoff vector.

I recommend checking your understanding of b and y , including how they are constrained in the construction of an equilibrium of the game from the current period. The first thing to note is that incentives in the production phase of the current period are influenced by both the current-period transfer b and the continuation value y . Second, y incorporates both self-enforcement and external enforcement of future behavior.

Third, verifiability constraints limit the scope of b , requiring this function to be P -measurable. Verifiability constraints apply also to the externally enforced aspect of future behavior, but the self-enforced aspect of future behavior (the manner in which the players coordinate future actions to achieve a continuation value) is not constrained by verifiability because the players commonly observe x . In this sense y is less constrained than is b . However, for any fixed x , whereas $b(x)$ is unbounded, $y(x)$ must be in the set of equilibrium values in the continuation of the game from the next period. Hence, a major theme of the analysis of enforcement is that there is a trade-off between current-period transfers and the next period's continuation value in the provision of incentives.

Although I have not yet provided details about the negotiation phase or specified an equilibrium notion, we can begin to characterize equilibrium behavior in the negotiation phase in terms of supportable continuation values. We have defined $y: X \rightarrow \mathbb{R}^n$ to give the continuation value from the start of the next period as a function of the current-period outcome x , and we know that y is constrained by enforcement and verifiability conditions in the continuation of the game. These constraints can be represented by a subset Y of functions from X to \mathbb{R}^n . Again, we haven't defined the full model yet, but it is enough to recognize that there will be *some* set of achievable functions that give the continuation value from the next period, and so let us call it Y .

Consider any P -measurable transfer function $b: X \rightarrow \mathbb{R}_0^n$, any continuation-value function $y \in Y$, and any mixed action profile $\alpha \in \Delta A$. Suppose the players enter the production phase with a contract that specifies b for the current period, continuation values in the next period given by y , and play of α in the current period. The key question is whether α is self-enforced, meaning that it is a Nash equilibrium of the induced game in Expression 3. If so, we know that, from the production phase in the current period, the following continuation value is achievable:

$$(1 - \delta) [u(\alpha) + \bar{b}(\alpha)] + \delta \bar{y}(\alpha).$$

If we add the transfer made in the negotiation phase of the current period, then we get the continuation value from the start of the current period.

The foregoing analysis provides a way to calculate the set of achievable continuation values from the production phase of a given period, as a function of the current-period transfer function b and the set Y of continuation-value functions. For any such b and Y , where b is P -measurable, let us define

$$D(b, Y) \equiv \left\{ (1 - \delta) [u(\alpha) + \bar{b}(\alpha)] + \delta \bar{y}(\alpha) \mid y \in Y \text{ and } \alpha \text{ is a Nash equilibrium of } \langle A, (1 - \delta)(u + \bar{b}) + \delta \bar{y} \rangle \right\}. \quad (4)$$

This operator is the core element of the recursive technique for characterizing equilibria of repeated-game and relational-contracting models. The operator is monotone in that $D(b, Y) \subset D(b, Y')$ for $Y \subset Y'$. For the rest of the analysis, we have to fill out the details of the model by describing exactly what happens in the negotiation phase of each period.

2.4 An aside: repeated games and the APS algorithm

Before analyzing versions of the relational-contracting model in detail, it is useful to consider a standard repeated game with imperfect public monitoring, where there is no negotiation phase and no contractible transfers. In each period, the players only choose actions in a stage game $\langle A, X, \lambda, u \rangle$ and receive payoffs. The components of the stage game are the same as in the relational-contracting model, but without the partition P because there is no external contract to study. The stage game is the same in every period.

Such a repeated game can be analyzed in the framework of our relational-contracting model by ignoring P , assuming that nothing takes place in the negotiation phase (no trans-

fers, in particular), and assuming that the externally enforced transfer function is exogenously fixed at $b \equiv \mathbf{0}$ in every period. Let us also assume that the players have access to an arbitrary public randomization device at the end of each period; this enables the players to jointly randomize over their future plans.

The standard solution concept for this repeated-game model is perfect public equilibrium (PPE), which is a strategy profile that satisfies two conditions. First, each player best responds in the continuation of the game from every period and for every history of play. Second, the strategies are functions only of the public history. That is, the equilibrium action profile in period t depends only on $\{x^1, x^2, \dots, x^{t-1}\}$.

There are generally many perfect public equilibria, and thus many different equilibrium payoff vectors. Abreu, Pearce, and Stacchetti (1990) established a relation between PPE and a set operator like that developed in the previous subsection. To be precise, for any set $W \subset \mathbb{R}^n$, let $F(W)$ to be the set of all functions from X to $\text{conv } W$, where “conv W ” denotes the convex hull of W :

$$F(W) \equiv \{y: X \rightarrow \text{conv } W\}.$$

The convex hull captures the effect of the public randomization device. Abreu, Pearce, and Stacchetti (1990) prove that the set of PPE payoff vectors in the repeated game is exactly the largest fixed point of the operator $D(\mathbf{0}, F(\cdot))$.

The recursive technique allows us to characterize PPE continuation values without describing full strategies. One procedure for calculating W^* takes advantage of its monotone property; we start with the convex hull of the set of all feasible payoff vectors and apply the operator $D(\mathbf{0}, F(\cdot))$ iteratively to find a limit set. This procedure is sometimes difficult to perform because it operates in the space of subsets of \mathbb{R}^n . But the counterpart for relational-contract settings is simplified by the special structure of these models, as shown below.

3 Relational Contracts with Trivial External Enforcement

In this section I describe the analysis of settings in which the externally enforced transfer function b is exogenously fixed and the same in every period. The interpretation is that either there is no external enforcement authority at all, or the enforcement technology requires commitment to a single transfer function over time. In the latter case, we could imagine that the players or an outside party selects the transfer function before the relationship starts. Let us continue to assume that the players have access to an arbitrary public randomization device at the end of each period.

3.1 Passive contracting

In the first version of the model, all that happens in the negotiation phase is that the players simultaneously make voluntary transfers, modeled noncooperatively. Each player can transfer any nonnegative amounts of money to the others and can also throw money away.

These transfers are observed by everyone. The vector sum of transfers defines $m \in \mathbb{R}_+^n$, the total transfer in the negotiation phase. Since there is no real negotiation accounted for in the negotiation phase, I call this a relational-contracting environment with “passive contracting.”¹² Quite a few relational-contract models essentially fall into this category.¹³

As with the class of repeated games described in Section 2.4, perfect public equilibrium is the solution concept typically used to analyze this model, and it can be expressed in terms of a recursive formulation of equilibrium continuation values from each phase of the game. Let W^* denote the set of PPE continuation values from the start (negotiation phase) of every period, as before. Additionally, let W' be the set of PPE continuation values from the production phase in each period. These two sets are of course related. Incentive conditions in the production phase require $W' = D(b, F(W^*))$ because, after production, interaction continues in the negotiation phase of the next period. Likewise, incentive conditions for the negotiation phase are captured by an operator that makes W^* a function of W' .

To define the second operator, let us work out what continuation values can be supported from the negotiation phase when the players must coordinate on values in a given set W from the production phase of the current period. The players choose voluntary transfers, resulting in the sum transfer m . Then, as a function of the transfers made, they proceed to coordinate on a particular continuation value $w \in W$. Normalizing the transfer, this produces a continuation value of $(1 - \delta)m + w$ from the negotiation phase.

Note that the most severe way to punish player i for not making the required transfer would be to coordinate on a continuation value in W that minimizes w_i . Player i can guarantee herself at least this amount from the negotiation phase by transferring nothing to the others. It turns out that any feasible continuation value that gives each player at least his or her minimum level can be achieved in a way that satisfies all of their incentive conditions in the negotiation phase.¹⁴ Allowing for the fact that the minimum I just referred to may not exist, so the infimum is needed, define

$$\text{tri } W \equiv \left\{ (1 - \delta)m + w \mid m \in \mathbb{R}_+^n, w \in W, \text{ and for every } i, \right. \\ \left. \text{there exists } \underline{w}^i \in W \text{ such that } (1 - \delta)m_i + w_i \geq \underline{w}_i^i \right\}.$$

Here *tri* stands for “triangle;” the reason for this name will be apparent shortly. The sets of PPE continuation values must satisfy $W^* = \text{tri } W'$.

The PPE characterization follows by combining operators D and *tri*. Because these operators are both monotone, the composition $\text{tri } D(b, F(\cdot))$ is also monotone. We have as a result that the set of PPE payoff vectors in the relational-contract game is exactly the largest fixed point of the operator $\text{tri } D(b, F(\cdot))$, which we have denoted W^* .

Here is where transferable utility, the simplifying assumption made by MacLeod and Malcomson (1989) and many others since, starts to deliver benefits in the characterization

¹²See Watson (2013) for a discussion of how to model variations in the “activeness of contracting.”

¹³These include, for example, MacLeod and Malcomson (1989) and Levin (2003) on principal-agent relationships, Doornik (2006) and Schöttner (2008) on partnerships/team production. They include extra phases and (except for Schöttner 2008) separation choices in each period. The threat of separation imposes a lower bound on equilibrium continuation values, but otherwise, there is not much of consequence that differs.

¹⁴This is not difficult to show. See Goldlücke and Kranz (2012, 2013) for general analysis.

of equilibrium. Observe that, whatever is W , $\text{tri } W$ is a generalized triangle with a linear frontier of slope -1 . If W is compact, there are numbers $\underline{w}_1^1, \underline{w}_2^2, \dots, \underline{w}_n^n$ such that, letting $L = \max_{w \in W} \sum_{i=1}^n w_i$ denote the *level*, we have $v \in \text{tri } W$ if and only if $\sum_{i=1}^n v_i \leq L$ and $v_i \geq \underline{w}_i^i$. Thus, the PPE value set W^* has the same characterization and we name its level L^* . Every element of W^* splits the joint value of L^* arbitrarily between the players, with free disposal and such that each player i gets at least her minimum \underline{w}_i^i .

In other words, a vector is in W^* if and only if it can be expressed as the players jointly getting level L^* and making a monetary transfer that is unrestricted except that each player must obtain at least her minimum continuation value. Thus, specifying a continuation value to be received at the start of the next period is just like picking a transfer to be received in the current period, factoring in discounting. The set W^* is characterized by $n + 1$ numbers and, for relatively simple production technologies, it becomes straightforward to calculate. Furthermore, we do so without having to describe the equilibrium strategies.

Perfect public equilibrium in the project-choice example

Consider the example described in Section 1.2, with parameters shown in Table 1 and stage-game payoffs pictured in Figure 1. In the current setting of trivial external enforcement, the externally enforced bonus payment ρ is constant across periods. We can characterize the PPE set for any fixed value of ρ and parameters δ and π , and then identify a value of ρ that maximizes the level L^* . Let us work through the analysis in the cases of $\rho = 4$ and $\rho = 28$. It turns out that the level is no higher for any other bonus. Since ρ defines b , given the normalization that the transfer is zero under the low signal, I write $D(\rho, F(W^*))$ rather than $D(b, F(W^*))$ in expressions to follow.

Start with the case of $\rho = 4$. From the characterization result described above, we know that the PPE value set is of the form

$$W^* = \text{conv}\{(\underline{w}_1^1, L^* - \underline{w}_1^1), (L^* - \underline{w}_2^2, \underline{w}_2^2), (\underline{w}_1^1, \underline{w}_2^2)\}.$$

In the production phase, by choosing $a_1 = 0$ or $a_1 = 2$, player 1 can guarantee himself a continuation value of at least $(1 - \delta) \cdot 0 + \delta \underline{w}_1^1$. Similarly, player 2's continuation value is bounded below by $(1 - \delta) \cdot 0 + \delta \underline{w}_2^2$. It is also easy to see that $a_1 = 0$ is a Nash equilibrium of the induced game with constant continuation-value function $y(0) = y(1) \equiv (\underline{w}_1^1, \underline{w}_2^2)$, which establishes that $\delta(\underline{w}_1^1, \underline{w}_2^2)$ is for both players the lowest value in the set $D(4, F(W^*))$ and hence the same in $\text{tri } D(4, F(W^*))$. Therefore $\delta(\underline{w}_1^1, \underline{w}_2^2) = (\underline{w}_1^1, \underline{w}_2^2)$, which implies $\underline{w}_1^1 = \underline{w}_2^2 = 0$.

To finish calculating the PPE value set, we need only determine L^* , which will depend on the discount factor. Consider a candidate value set W with level L . The highest continuation value for player 1 from the negotiation phase of any period is L and the lowest is 0. Note that $a_1 = 2$ is a Nash equilibrium of the induced game in the production phase for any constant continuation-value function, which achieves value $(0, 6)$, and therefore L is at least 6.

To have $L = 8$, player 1 must be given the incentive to select $a_1 = 1$. Presuming that this is the case, note that the greatest incentive for $a_1 = 1$ is provided as follows:

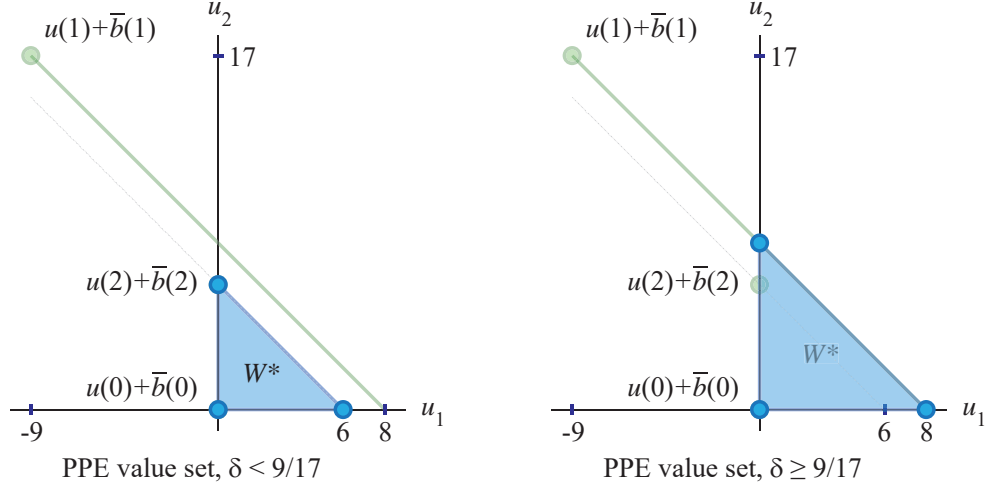


Figure 2: PPE values in the example with $\rho = 4$.

Contingent on $a_1 = 1$ in the current period, the players coordinate to achieve continuation value $(8, 0)$ from the negotiation phase of the next period; for any other a_1 (in particular for $a_1 = 0$ or $a_1 = 2$, which give the greatest deviation gain), they coordinate to achieve $(0, 8)$ from the next period. The resulting incentive condition is:

$$(1 - \delta)(-9) + \delta \cdot 8 \geq (1 - \delta) \cdot 0 + \delta \cdot 0,$$

which simplifies to $\delta \geq 9/17$.

Thus, if $\delta \geq 9/17$ then cooperation with the efficient project choice can be sustained and $L^* = 8$, and otherwise $L^* = 6$. Consistent with intuition central to repeated-game theory, cooperation can be sustained if the players are patient enough. The PPE value set W^* is pictured in Figure 2.

The analysis is similar in the case of $\rho = 28$. By choosing $a_1 = 2$ or $a_1 = 3$, player 1 can guarantee himself a continuation value of at least $(1 - \delta) \cdot 6 + \delta \underline{w}_1^1$ from the production phase. Similarly, player 2's continuation value is bounded below by $(1 - \delta) \cdot 0 + \delta \underline{w}_2^2$. Clearly $a_1 = 2$ is a Nash equilibrium of the induced game with constant continuation-value function $y(0) = y(1) \equiv (\underline{w}_1^1, \underline{w}_2^2)$, which establishes that $(1 - \delta)(6, 0) + \delta(\underline{w}_1^1, \underline{w}_2^2) = (\underline{w}_1^1, \underline{w}_2^2)$, implying $(\underline{w}_1^1, \underline{w}_2^2) = (6, 0)$. To reach level 8, player 1 must be given the incentive to select $a_1 = 1$ by rewarding with continuation value $(8, 0)$ and punishing with $(6, L - 6)$. The resulting incentive condition is:

$$(1 - \delta) \cdot 3 + \delta \cdot 8 \geq (1 - \delta) \cdot 6 + \delta \cdot 6,$$

which simplifies to $\delta \geq 3/5$. Thus, under this condition cooperation with the efficient project choice can be sustained and $L^* = 8$; otherwise $L^* = 6$.

Because the PPE level is at least 6 with both $\rho = 4$ and $\rho = 28$, and because the cutoff discount factor for level 8 is lower under $\rho = 4$, it is optimal to specify bonus $\rho = 4$ for the relationship, regardless of the discount factor. Cooperation can be sustained at the efficient project choice in the case of $\delta \geq 9/17$.

In their analysis of principal-agent settings, MacLeod and Malcomson (1989) observed that the worker’s incentive condition in the production phase and the manager’s incentive condition regarding compensating the worker in the negotiation phase can be pooled to form a single necessary and sufficient inequality. Levin (2003) further observed that, if cooperation can be sustained in a PPE, then it can be sustained in a *Pareto-perfect* PPE, where every equilibrium continuation value is on the efficient frontier of W^* . Goldlücke and Kranz (2013) provide general results for all settings with two players and perfect monitoring.¹⁵ But note also that, regardless of δ , inefficient PPE exist, such as the strategy that has the players never making transfers and player 1 choosing $a_1 = 0$ in the project-choice example with $\rho = 4$.

The nature of contract

Since there is only self-enforcement in this version of the model, the relational “contract” here is just whatever perfect public equilibrium the players have coordinated on. Because this model does not have active contracting, there is no theory of how the equilibrium is selected, much less that the players would coordinate on an equilibrium at all. In other words, the current model incorporates the self-enforced aspects of contracting but says nothing about the establishment of contracts. For a more complete model of relational contracting, we need to account for the contracting process, and this is where the next section heads.

3.2 Active contracting

The next step in our tour is to enrich the model with an explicit account of active contracting, where players exercise bargaining power in the process of reaching agreements. Whereas in the previous version of the model we assumed that the negotiation phase is just a time when players make voluntary transfers, we now assume that interaction in the negotiation phase includes actual bargaining over a contract and an immediate transfer. In the present context, where there is trivial external enforcement, a contract is an agreement only about future behavior to be self-enforced. As in the model of a short-term relationship, one can account for negotiation either noncooperatively or cooperatively. The former approach specifies a noncooperative bargaining protocol, where players make and respond to offers of contracts and immediate transfers. The latter approach specifies a cooperative bargaining solution to account for play in the negotiation phase.

¹⁵Here Pareto-perfection is imposed in the negotiation phase. Imposing the condition in a standard repeated game with voluntary transfers incorporated into the stage game, as examined by Baliga and Evans (2000), yields a less tractable characterization of equilibrium values. Pareto-perfection underlies definitions of *renegotiation-proof* equilibrium in the repeated-game literature, specifically those of Rubinstein (1980), Bernheim and Ray (1989), and Farrell and Maskin (1989). Pearce’s (1987) definition has a different foundation and allows for Pareto-ranked equilibrium continuation values. None of these theories of renegotiation-proofness actually model the negotiation process; they do not contemplate the possibility of disagreement, and bargaining power plays no role. In contrast, the theory of *contractual equilibrium*, discussed in the next subsection, is based on an explicit account of the negotiation process.

Miller and Watson (2013) and Watson (2013) introduced a framework for modeling relational contracts with active contracting, and this is what I'll focus on here.¹⁶ Miller and Watson (2013) develop a fully noncooperative model as well as a hybrid version in which the negotiation is modeled cooperatively using the Nash bargaining solution with fixed bargaining weights.¹⁷ They provide an equivalence result that connects the two approaches. The result is a theory of behavior called *contractual equilibrium* in both the fully noncooperative game and the hybrid game. Conveniently, the set of contractual equilibrium values has a recursive characterization along the lines of that for PPE values. So we do not need to fuss with a description of equilibrium strategies here, but instead describe the hybrid version of the model and the recursive formulation of equilibrium values.

We can think of the players, in the negotiation phase of any period, as bargaining over (i) an immediate transfer, (ii) the action profile they will play in the current period, and (iii) their coordinated behavior in future periods. The third element is summarized by their continuation value as a function of the current-period outcome. The continuation value incorporates the players' anticipated renegotiation of their agreement in future periods. One way to parse this is that the players agree how they will play in future periods unless and until they successfully renegotiate, so the value of the agreed behavior becomes the disagreement point for future renegotiation. They are thus implicitly bargaining over the implied continuation values that incorporate anticipated renegotiation in future periods.

Contractual equilibrium values

In the hybrid model, the bargaining protocol is represented by an exogenously fixed vector of bargaining weights $\pi = (\pi_1, \pi_2, \dots, \pi_n)$ satisfying $\pi_i \geq 0$ and $\sum_{i=1}^n \pi_i = 1$, applying to every period. Bargaining in the negotiation phase is resolved according to the generalized Nash bargaining solution, where the players coordinate to achieve a continuation value that maximizes their joint value and, by making an immediate transfer, distribute the surplus in proportion to their bargaining weights. The surplus is relative to a *disagreement point* that is assumed to entail no immediate transfer and then coordination to achieve an available continuation value from the production phase. The disagreement point may depend on the history of interaction to the previous period, implying that generally multiple continuation values can be supported from the negotiation phase.

Here is the characterization of the contractual equilibrium value set. Suppose a set W gives the continuation values from the negotiation phase of any period. Just as in the previous version of the model, incentive conditions in the production phase imply that $D(\mathbf{0}, F(W))$ is the set of equilibrium values from the production phase, and the maximum

¹⁶The theory builds on the hybrid modeling approach described in Watson (2001).

¹⁷Because players can threaten to "hold-up" negotiation, they take advantage of bargaining power in much the same way as in models of short-term trading relationships in which unverifiable investments are followed by verifiable trade, as in Williamson (1985) and Grossman and Hart (1986). For general analysis based on mechanism design, see Maskin and Moore (1999) and Watson (2007).

joint value is

$$L = \max_{w \in D(\mathbf{0}, F(W))} \sum_{i=1}^n w_i .$$

The bargaining solution requires that every value $w \in W$ must satisfy

$$w = \underline{w} + \pi \left(L - \sum_{i=1}^n \underline{w}_i \right), \quad (5)$$

for some $\underline{w} \in D(\mathbf{0}, F(W))$. In this expression, \underline{w} is the disagreement point and the term in parentheses is the bargaining surplus. We say the set W is *bargaining self-generating* (BSG) if it satisfies this condition, and we call L its *level*. Clearly any BSG set has a constant joint value in that $\sum_{i=1}^n w_i = L$ for every $w \in W$. The contractual equilibrium value (CEV) set W^* is defined as the dominant BSG set in the sense of maximizing the level, and we let L^* denote its level.¹⁸ Under suitable technical conditions, contractual equilibrium exists and W^* is compact.

Figure 3 illustrates the CEV set in the two-player setting. Here the CEV set is a compact subset of a line of slope -1 , characterized by its endpoints z^1 and z^2 , where z^1 is the value that is worst for player 1 and z^2 is the worst point for player 2. In fact, whether any particular points in the interior of this line segment are included in W^* is inconsequential to the equilibrium construction because all such points can be achieved in expectation using the public randomization device. Therefore, we only need to keep track of the endpoints z^1 and z^2 . The *span* of W^* , denoted by d^* , is defined as the horizontal (equivalently vertical) length of the CEV set; that is, $d^* = z_1^2 - z_1^1 = z_2^1 - z_2^2$. The span figures prominently in the analysis of examples.

We can find the CEV set W^* by deconstructing the two endpoints. Associated with each endpoint z^i is a disagreement point $\underline{w}^i \in D(\mathbf{0}, F(W^*))$ such that the following holds:

$$z^i = \underline{w}^i + \pi \left(L^* - \underline{w}_1^i - \underline{w}_2^i \right) = \pi L^* + \left(\pi_2 \underline{w}_1^i - \pi_1 \underline{w}_2^i, \pi_1 \underline{w}_2^i - \pi_2 \underline{w}_1^i \right). \quad (6)$$

This is the condition of Equation 5. Because z^1 is the point in W^* that minimizes player 1's payoff, the associated disagreement point \underline{w}^1 is the point in $D(\mathbf{0}, F(W^*))$ that is furthest in the direction $(-\pi_2, \pi_1)$, orthogonal to π . Likewise, disagreement point \underline{w}^2 is the point in $D(\mathbf{0}, F(W^*))$ that is furthest in the direction $(\pi_2, -\pi_1)$. This is illustrated in the right graph of Figure 3.

Let us pause for a moment to interpret the solution concept and reflect on how features of the model relate to the characterization of the CEV set. Recall that, in the setting of passive contracting, transferrable utility and the ability to make transfers in the negotiation phase implies that the PPE value set is a generalized triangle and therefore characterized by $n + 1$ numbers. In the current setting of active contracting, the model's bargaining

¹⁸The dominance condition is straightforward to assess because BSG sets have constant joint values and are therefore ranked by level. It turns out that this ranking is also in terms of Pareto dominance, in the sense that if BSG sets W and W' have levels satisfying $L \leq L'$, then there exists a BSG set W'' with level $L'' \geq L'$ such that every point in $\text{conv } W$ is weakly Pareto dominated by a point in $\text{conv } W''$.

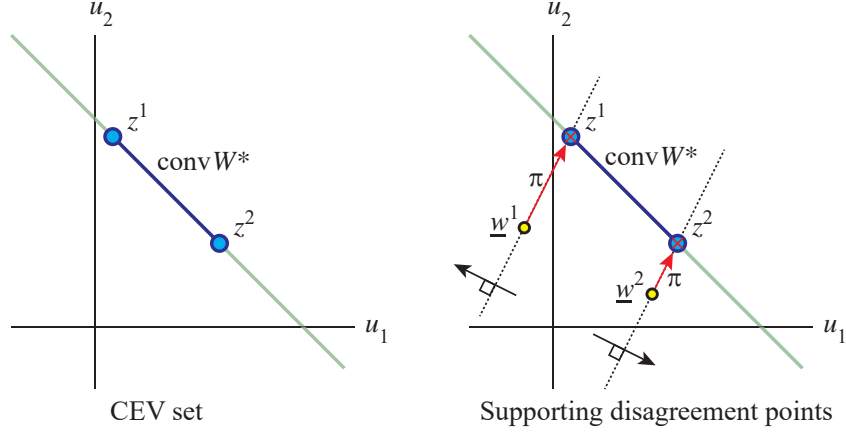


Figure 3: An illustration of the CEV set.

component implies that continuation values from the beginning of periods all have the same joint value, which puts more structure on the CEV set. For more intuition, in particular on steps to calculate the CEV set, let's return to the example.

Contractual equilibrium in the project-choice example

Similar to our analysis of PPE values, we can characterize the CEV set in our running example for any fixed value of ρ and parameters δ and π . Then we can identify a value of ρ that maximizes the level L^* . As before, it suffices to focus on the cases of $\rho = 4$ and $\rho = 28$, because the players can do no better with any other external contract.

Consider the case of $\rho = 4$ and let us take any candidate value set W with endpoints z^1 and z^2 , level L , and span d . It is not difficult to confirm that the point in $D(4, F(W))$ that is furthest in the direction $(\pi_2, -\pi_1)$ is achieved by specifying $y(0) = y(1) = z^2$ and having player 1 choose $a_1 = 0$, which is an equilibrium of the induced game. Player 1's payoff is maximized and player 2's payoff is minimized. Therefore $\underline{w}^2 = (1 - \delta)(0, 0) + \delta z^2$. Using Equation 6 for $i = 2$, substituting for \underline{w}^2 , using the fact that $z_1^2 + z_2^2 = L^*$, and rearranging terms yields $z^2 = \pi L$.

We can also confirm that, if it is possible to give player 1 the incentive to choose $a_1 = 1$, then the point in $D(4, F(W))$ that is furthest in the direction $(-\pi_2, \pi_1)$ is achieved by specifying $y(1) = z^1 + (9, -9)(1 - \delta)/\delta$, $y(0) = z^1$, and play of $a_1 = 1$. We obtain this specification of y by combining the incentive condition for $a_1 = 1$, namely

$$(1 - \delta)(-9) + \delta y_1(1) \geq (1 - \delta) \cdot 0 + \delta y_1(0),$$

with the objective of reducing player 1's payoff by making $y(1)$ as small as possible while also keeping $y(0)$ in W . We then have

$$\underline{w}^1 = (1 - \delta)(-9, 17) + \delta \left(z^1 + (9, -9) \cdot \frac{1 - \delta}{\delta} \right) = (1 - \delta)(0, 8) + \delta z^1.$$




Figure 4: CEV set in the example with $\rho = 4$.

Achieving $a_1 = 1$ implies $L = 8$, and so $z_1^1 + z_2^1 = 8$ and $\underline{w}_1^1 + \underline{w}_2^1 = 8$. The surplus of negotiation in this case is zero, and Equation 6 for $i = 1$ then amounts to $z^1 = \underline{w}^1 = (1 - \delta)(0, 8) + \delta z^1$, which simplifies to $z^1 = (0, 8)$.

But note that this construction relies on $y(1)$ being an element of W , which requires $d \geq 9(1 - \delta)/\delta$. From our analysis of the other endpoint, we have $z^2 = 8\pi$, implying $d = 8\pi_1$. Thus, $a_1 = 1$ is indeed viable only if $8\pi_1 \geq 9(1 - \delta)/\delta$, which simplifies to $\delta \geq 9/(9 + 8\pi_1)$.

In conclusion, if $\delta \geq 9/(9 + 8\pi_1)$ then cooperation with high effort can be sustained, $L^* = 8$, $d^* = 8\pi_1$, and the CEV set W^* has endpoints $z^1 = (0, 8)$ and $z^2 = 8\pi$. Otherwise, in equilibrium player 1 must select $a_1 = 0$ or $a_1 = 2$ in the production phase, $L^* = 6$, $d^* = 6\pi_1$, and W^* has endpoints $z^1 = (0, 6)$ and $z^2 = 6\pi$. The CEV set is shown in Figure 4.

The case of $\rho = 28$ is similarly analyzed. If $\delta \geq 3/(3 + 2\pi_1)$ then cooperation with high effort can be sustained, $L^* = 8$, $d^* = 2\pi_1$, and W^* has endpoints $z^1 = (6, 2)$ and $z^2 = (6, 0) + 2\pi$. Otherwise, in equilibrium it is only possible to have $a_1 = 2$ or $a_1 = 3$ in the production phase, $L^* = 6$, and W^* is the singleton with $z^1 = z^2 = (6, 0)$. Observe that the level is at least 6 with both $\rho = 4$ and $\rho = 28$; further, the cutoff discount factor for level 8 is lower under $\rho = 4$. Therefore $\rho = 4$ is optimal for the relationship, regardless of the discount factor. Notice that, whereas $\rho = 28$ minimizes player 1's gain of deviating from $a_1 = 1$ in the stage game, the relatively small span makes this external contract inferior to $\rho = 4$.

Comparing the results here to those in the previous subsection, the implications of active contracting are clear. First, because players can renegotiate in every period and would pursue their mutual interest in increasing their joint value, in equilibrium they are always at the frontier of their achievable set of continuation values from the negotiation phase. Thus, whereas the PPE value set in the setting of passive contracting generally has

Pareto-ranked elements, the CEV set in the setting of active contracting does not.

Second, the players' bargaining weights affect the CEV set. In this example, disagreement point \underline{w}^2 entails selection of a project in the current period that is inefficient compared to the choice that defines \underline{w}^1 . The renegotiated continuation value z^2 adds a surplus, of which player 1 obtains fraction π_1 . Indeed, this is the continuation value that the players would coordinate on *to reward player 1 for choosing the efficient project in the previous period*, so providing player 1 with incentives requires that he have sufficient bargaining power to extract a substantial share of the renegotiation surplus. In other words, when it comes time for player 2 to compensate player 1 for exerting high effort in the previous period, player 2 can hold up player 1 and try to avoid the payment. But this would lead player 1 to exert low effort in the current period, so they renegotiate and player 1 uses her bargaining power to extract the transfer from player 2.

3.3 Variations and extensions

Miller and Watson's (2013) analysis is based on assumptions about the meaning of statements that the parties make to each other. Alternative assumptions lead to variations of the contractual-equilibrium concept. For instance, Ramey and Watson (1999, 2002) and Klimenko, Ramey, and Watson (2008) invoke a tighter theory of bargaining in which the disagreement point in any given period cannot depend on the history of interaction. An implication is that contract enforcement requires group punishments that enforcement institutions may be needed to facilitate.¹⁹ Goldlücke and Kranz (2019) develop a similar model in which renegotiation is triggered randomly, and when it occurs, the disagreement point depends only on a state verifiable representing a prior technology choice. Safronov and Strulovici (2018) adopt a more permissive theory of bargaining than do Miller and Watson (2013), in which players may be punished for proposing Pareto improvements; this leads to a wider range of equilibrium outcomes than predicted by contractual equilibrium.

As mentioned already, a number of relational-contracting models include an option for unilateral separation. Models in the macro-labor literature typically combine relational contracting between workers and firms with matching markets; thus, when the relationship between a worker and firm is severed, the worker enters the unemployment pool to await a match with another firm, while the firm enters the pool of employers with vacant positions. Relational contracts interact with conditions in the matching market and in other markets.²⁰ The effect on attainable values of the option to separate depends on whether separation is assumed to be triggered by disagreement in the negotiation phase, where it becomes the disagreement point, or is an option at another time in each period, where it may be chosen to punish a deviation. Generally, increasing the value of the players' outside options tightens incentive constraints within a relationship.

¹⁹This may be particularly salient in international trade. Samples of recent theoretical studies are Beshkar (2016) and Buzard (2017).

²⁰Articles in this category include Shapiro and Stiglitz (1984), MacLeod and Malcolmson (1989), Ramey and Watson (1997), and den Haan, Ramey, and Watson (2000), and Sobel (2006).

Another active direction of research is analysis of relational contracts with private information. In one category are models with imperfect private monitoring, such as where a manager receives a private signal of the worker’s effort. If we assume that the manager can make a statement about her signal, which becomes part of the outcome of production, then this setting fits into the framework described here. A second category comprises models with persistent private information, such as if a worker’s effort cost is determined by nature at the beginning of the game and known only to the worker, which are well outside our framework.²¹

4 Relational Contracts with Nontrivial External Enforcement

Let us next examine settings with nontrivial external enforcement. The literature contains a variety of modeling exercises in which self-enforced and externally enforced terms interact, but most studies substantially limit the extent of external enforcement and/or make ad hoc assumptions about equilibrium selection. Here I’ll sketch a special case of the general model of Watson, Miller, and Olsen (2020), which does not have such limitations and has the added benefit of including a recursive formulation of equilibrium values. It is a setting of active contracting and the solution concept is contractual equilibrium. At the end of this section I briefly describe some of the other theories in the literature.

Let us drop the assumption made in Section 3 that the transfer function b is exogenously fixed at $\mathbf{0}$ for all periods, so we return to the full generality of the model sketched in Section 2. I now provide more details. Let $\mathcal{B} \equiv \{b: X \rightarrow \mathbb{R}_0^n\}$ denote the set of transfer functions, and let $H^X \equiv \cup_{k=0}^{\infty} X^k$ be the space of finite-length outcome histories, where the element for $k = 0$ is defined as the null history h^0 at the start of the relationship. An external contract specifies a transfer function b^t for each period t , itself as a function of the history of outcomes through period $t - 1$. To be formal, an external contract is a function $c: H^X \rightarrow \mathcal{B}$, where for any $(t - 1)$ -period history $h \in H^X$, the transfer function specified for period t is $b^t = c(h)$.

It is most convenient to deal with these contracts in the form of “continuation contracts.” Given a history of outcomes through period $t - 1$, the continuation contract from period t gives b^τ in each period $\tau \geq t$ as a function of the history of outcomes from t until $\tau - 1$. In other words, for any fixed history to period $t - 1$, we keep track of the contracted transfer functions starting from period t , as a function of the outcomes from period t on.

The continuation contract in a given period may be interpreted as specifying the transfer function b for the current period and a mapping from current-period outcome x to the continuation contract in force at the beginning of the next period. Formally, for any $x \in X$ and $h \in H^X$, where h is k periods in length, let xh denote the $(k + 1)$ -period outcome

²¹MacLeod (2003), Levin (2003), and Fuchs (2007) analyzes settings with private monitoring. Persistent private information is found in the models of Watson (1999, 2002), Halac (2012), Malcomson (2016), and Fahn and Klein (2019).

history in which x is followed by the sequence h . Define $c|x: H^X \rightarrow \mathcal{B}$ by $(c|x)(h) \equiv c(xh)$ for every $h \in H^X$. If the players operate under continuation contract c in period t , then they have transfer function $c(h^0)$ in this period and, after realizing outcome x , they will enter the following period with continuation contract $c|x$.

Because external contracts can depend only on information that is verifiable, 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. Let C be the set of contracts that respect verifiability.²²

4.1 Active contracting

In the setting of nontrivial external enforcement, contracting cannot be completely passive because a negotiation protocol is required to model how the players can make changes to the external contract. While it is possible to model contracting that is passive to some degree, I shall skip such an exercise and proceed directly to active contracting along the lines of the model described in Section 3.2. Let us focus on the case of two players, where $n = 2$.

Players begin their relationship in period 1 with the default external contract \hat{c}^1 that specifies transfer function $\mathbf{0}$ for every period regardless of the history. Then in the negotiation phase of the first period, the players bargain over a new external contract c^1 to replace \hat{c}^1 , an immediate transfer m^1 , and a specification of future behavior summarized by continuation values. The disagreement point entails $c^1 = \hat{c}^1$ and $m^1 = (0, 0)$. At the end of the first period, the productive outcome x^1 determines the continuation contract $\hat{c}^2 = c^1|x^1$ that the players inherit at the start of period 2.

The negotiation phase works the same way in every future period t . The players enter the period with continuation contract \hat{c}^t , they negotiate to change it to c^t and make transfer m^t , and then the outcome x^t determines $\hat{c}^{t+1} = c^t|x^t$ for period $t + 1$. The disagreement point entails $c^t = \hat{c}^t$ and $m^t = (0, 0)$. Note that the players bargain over both externally enforced and self-enforced components of their contract in the negotiation phase, so there is more happening in this model than we had in the setting of Section 3.2.

Contractual equilibrium

Let us account for interaction in the negotiation phase cooperatively, where 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 $\pi_1 + \pi_2 = 1$. Contractual equilibrium can be represented by a recursive formulation of continuation values as before. Because external contracts render the relational-contracting game nonstationary, the set of continuation values attainable from a given period depends on the inherited contract. Let $W(c) \subset \mathbb{R}^2$ denote the set of continuation values from the beginning of a period in which c is the inherited contract, and let $\mathcal{W} = \{W(c)\}_{c \in C}$ be the collection.

²²A contract c respects verifiability if, for all $x, x' \in X$, $x \in P(x')$ implies $c|x = c|x'$.

Let us first consider what continuation values can be achieved from the production phase of a period under continuation contract c . Note that for any outcome x , the contract inherited in the next period will be $c|x$ and so the continuation value from the start of the next period must be in the set $W(c|x)$. This means that the set of feasible continuation-value functions is

$$F^c(\mathcal{W}) \equiv \{y: X \rightarrow \mathbb{R}^2 \mid y(x) \in \text{conv } W(c|x) \text{ for every } x \in X\}.$$

Recalling that c specifies transfer function $b = c(h^0)$ in the current period, we find that the set of continuation values attainable from the production phase is $D(c(h^0), F^c(\mathcal{W}))$.

Next we apply the bargaining solution. In the negotiation phase under inherited contract \hat{c} the players would coordinate on some value $\underline{w} \in D(\hat{c}(h^0), F^{\hat{c}}(\mathcal{W}))$ in the event that they fail to make an agreement, making \underline{w} the disagreement point for negotiation. The Nash bargaining solution predicts that the players renegotiate to a contract c and coordinate on a continuation value that maximizes their joint value,

$$L \equiv \max\{w_1 + w_2 \mid c \in C, w \in D(c(h^0), F^c(\mathcal{W}))\}, \quad (7)$$

and they choose an immediate transfer to achieve continuation value

$$w = \underline{w} + \pi(L - \underline{w}_1 - \underline{w}_2). \quad (8)$$

A collection $\mathcal{W} = \{W(c)\}_{c \in C}$ is called *bargaining self-generating* (BSG) if for every $\hat{c} \in C$ and $w \in W(\hat{c})$, there exists a value $\underline{w} \in D(\hat{c}(h^0), F^{\hat{c}}(\mathcal{W}))$ such that Equation 8 holds. We call L the *level* of the collection. Then the contractual equilibrium value (CEV) collection $\mathcal{W}^* = \{W^*(c)\}_{c \in C}$ is defined as the dominant BSG collection in the sense of maximizing the level, denoted L^* . Under suitable technical conditions, contractual equilibrium exists.

This model may seem impossible to solve, because the set of external contracts is huge and it is not obvious how to even begin the analysis of any example. Contracted transfers may depend on the outcome history in a nonstationary manner. For instance, in a principal-agent setting, a contract could specify a schedule of bonus payments that changes in response to past outcomes, ratcheting up or down over time. Several questions must be raised. What are the properties of the optimal external contract? Do the players renegotiate it on or off the equilibrium path? Does the external enforcement technology complement self-enforcement?

Fortunately, Watson, Miller, and Olsen (2020) provide a characterization result that applies to the model sketched here (their model is more general), greatly simplifying the analysis and helping to answer the questions just now posed. The optimal continuation contract c^* , which the players select to achieve the level L^* , is *semistationary*. It specifies one transfer function b^* for the first period and another transfer function \underline{b} for all other periods. There is no dependence on the history of outcomes. Thus, in equilibrium, in the first period the players agree to the external contract that specifies b^* for period 1 and \underline{b} for every period 2, 3, . . . , regardless of the history.

Furthermore, both on and off the equilibrium path, in each period the players renegotiate back to this same continuation contract. That is, in period 2 they revise the external contract to specify b^* in period 2, retaining the specification of \underline{b} for all future periods. In period 3 they revise again to specify b^* for period 3, and so on. The transfers they make in the renegotiated deals depend on the history because the manner in which they coordinate in disagreement depends on past outcomes.

The intuition behind this result has two parts: First, transfers specified in b can substitute for variations in the continuation contract, because they are conditioned on the same information, and this substitution can be done while preserving any needed variations in the self-enforced aspects of continuation value. This means that the continuation contract can be constant in the outcome of the current period. Second, what matters for incentives in the current period is the span of continuation values. By specifying a transfer function for all future periods to achieve the largest span, the players will be able to achieve the highest attainable joint value in the current-period stage game, and future renegotiation will ensure that the high value is achieved in future periods as well, without reducing the span of continuation values.

On the technical side, the characterization result provides an algorithm for calculating b^* and \underline{b} . The latter is determined as follows: Suppose we exogenously fix a single transfer function \hat{b} for all periods as in Section 3, a stationary setting with trivial external enforcement. Then we can calculate the game's CEV set—call it $\hat{W}(\hat{b})$ —and see how it depends on \hat{b} . It turns out that \underline{b} is the transfer function that maximizes the span of $\hat{W}(\hat{b})$. Then b^* is the transfer function that maximizes the players' joint value in the induced game in which all continuation values are in $\hat{W}(\underline{b})$.

Kostadinov (2020), in work contemporaneous with Watson, Miller, and Olsen (2020), proved a similar result for a principal-agent setting with risk aversion, utilizing the PPE solution concept. In Kostadinov's model, the parties form a semistationary external contract specifying one bonus scheme for the first period and a second bonus scheme for all future periods. Then in every period on or off the equilibrium path, the parties revise the contract to provide the former bonus scheme in the current period.²³

Contractual equilibrium in the project-choice example

Our running example illustrates contractual equilibrium in the setting with nontrivial external enforcement, as well as attendant economic insights regarding the functionality of nonstationary contracts. As before, please refer to Table 1 and Figure 1 for a reminder of the parameters. The analysis of contractual equilibrium in the setting of trivial external enforcement provides guidance for determining both \underline{b} and b^* .

Let us first put ourselves in the stationary setting of Section 3, where the bonus is fixed at ρ for all periods. Recall that in Section 3.2 we reached the conclusion that bonus $\rho = 4$ is optimal because it leads to a lower cutoff discount factor for the efficient project choice

²³Other reasons for optimal contracts to be nonstationary in time-invariant environments are 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 2016).

than would bonus $\rho = 28$. The incentive condition for $a_1 = 1$ resulted from comparing the deviation gain, which is $9(1 - \delta)$ in the case of $\rho = 4$ and $3(1 - \delta)$ in the case of $\rho = 28$, with the span used for rewards and punishments from the next period. Under parameter values that yield $L = 8$ where $a_1 = 1$ is achieved, the span of the CEV set is $8\pi_1$ in the case of $\rho = 4$ and $2\pi_1$ in the case of $\rho = 28$.

With nontrivial external enforcement, the players can specify different bonuses for the current and future periods, effectively coupling a small deviation gain in the current period with a large span from the next period. Importantly, bonus $\rho = 4$ results in a CEV span of at least $6\pi_1$ even if it is not possible to achieve the efficient project choice $a_1 = 1$ under this bonus. Further, bonus $\rho = 28$ minimizes the current-period deviation gain. Thus, by specifying $\rho = 28$ for the current period and $\rho = 4$ for all future periods, player 1 is given the greatest possible incentive to choose $a_1 = 1$ in the current period. The incentive condition is $3(1 - \delta) \geq L^*\pi_1\delta$, which simplifies to $\delta \geq 3/(3 + L^*\pi_1)$.

The conclusion is thus: If $3/(3 + 6\pi_1) \leq \delta \leq 9/(9 + 8\pi_1)$, then efficient effort can be sustained, $L^* = 8$, and the optimal external contract is semistationary. The contract specifying a large bonus of $\rho = 28$ in the first period and a smaller bonus $\rho = 4$ in all future periods is optimal. In each period, the players revise the bonus scheme to increase the current-period bonus to $\rho = 28$; renegotiation does not change the implied span of continuation values, which is $6\pi_1$. If $\delta \geq 9/(9 + 8\pi_1)$ then efficient effort can be sustained with a stationary external contract that provides for a bonus of $\rho = 4$ in every period. If $\delta < 3/(3 + 6\pi_1)$ then efficient effort cannot be sustained in equilibrium.

Notably we have found that, with nontrivial external enforcement, a semistationary contract delivers the efficient outcome in settings where efficiency is not attainable with a constant bonus (as studied in Section 3.2). This is the case for δ between $3/(3 + 6\pi_1)$ and $9/(9 + 8\pi_1)$. Thus, for applications, it is important to carefully consider nonstationary contracts and renegotiation.

4.2 Variations and extensions

Prior to Watson, Miller, and Olsen (2020) and Kostodinov (2019), most models of relational contracting with negotiation and nontrivial external enforcement restricted attention to short-term external contracts, as in Radner (1985) and Pearce and Stacchetti (1998), or stationary long-term external contracts, as in Che and Yoo (2001) and Itoh and Morita (2015). Prior theories are also varied in terms of whether and when active negotiation is assumed to occur, and whether players are able to renegotiate over one or both parts of their contract. For instance, Baker, Gibbons, and Murphy (1994, 2002) and Schmidt and Schnitzer (1995) assumed any deviation triggers an end to intertemporal self-enforcement, meaning that play in each future period must be a Nash equilibrium of the induced game with constant continuation values. But they also assumed that, following a deviation, the players would be able to renegotiate the external contract.²⁴ An implication is that im-

²⁴A similar line is taken by Kvaløy and Olsen (2009) and Iossa and Spagnolo (2011). And plenty of models with external enforcement assume that contracts are formed in a fairly passive way (such as via a

proving the external enforcement technology can have the effect of tightening incentive conditions and reducing welfare. In contrast, Watson, Miller, and Olsen (2020) show that, in the more realistic setting in which players can renegotiate both components of their contract, the external-enforcement technology always complements self-enforcement.

5 Directions for Further Study

The modeling framework presented here incorporates the essential elements of relational incentive contracts, is straightforward to work with, and suits a variety of applications. This simple and general platform is a good starting point for a research projects in a variety of areas. One promising line of research is to broaden the theoretical foundations of active contracting and bargaining power, motivated by questions such as the following. What might make bargaining over an external contract different from bargaining over the internal, self-enforced part? On either dimension, are parties more easily able to negotiate over some aspects of their relationship relative to other aspects? How do they establish or lose trust with respect to the connection between what they pledge to do and actually do? What happens if trust and coordination break down completely? These questions require a deeper analysis of bargaining and equilibrium selection than has been summarized in this article, and may also require theories of chaos.

There are technological variations to explore, especially for contractual equilibrium and related concepts. I have in mind issues such as costly contracting, bounds on transfers, nonstationary settings with state variables (such as debt holdings or capital), endogenous monitoring and technology choice, long-term investments, and short-term shocks.²⁵ Another research area ripe for further investigation is contracting in settings with overlapping relationships, both in time such, as with overlapping generations, and in space, such as with networks or communities.²⁶ Finally, there is plenty to explore in settings with persistent private information. In all of these areas, contributions that incorporate active contracting and deepen the foundations of bargaining would be well received.

Nash-demand protocol as in Rayo 2007) or simply do not allow for renegotiation (Barron et. al. 2019 is an example).

²⁵Existing models of long-term investments and/or shocks include Ramey and Watson (1997), Li, Matouschek, and Powell (2017), and Englmaier and Fahn (2019). Contracting costs have been studied outside of the relational-contract setting by Dye (1985), Battigalli and Maggi (2002, 2008), and Schwartz and Watson (2004). A state variables is present in Barron, Li, and Zator's (2019) model of a relational contract with debt. Modeling exercises with endogenous monitoring were done by Fong and Li (2016) and Gjertsen et. al. (2021).

²⁶Models of supplier management have been developed by Barron and Andrews (2016) and Board (2011). Ghosh and Ray (1996), Fujiwara-Greve and Okuno-Fujiwara (2009), Ali and Miller (2013, 2016), and Ali, Miller, and Yang (2017) construct models of relational contracting in communities and networks.

References

- Abreu, Dilip, Paul Milgrom, and David Pearce, "Information and Timing in Repeated Partnerships." *Econometrica* 59 no. 6 (1991): 1713–1733.
- Abreu, Dilip, David G. Pearce, and Ennio Stacchetti, "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* 58 no. 5 (1990): 1041–63.
- Ali, S. Nageeb and David A. Miller, "Ostracism and Forgiveness," *American Economic Review* 106, no. 8 (2016): 2329–48.
- Ali, S. Nageeb and David A. Miller, "Enforcing Cooperation in Networked Societies," unpublished manuscript (2013).
- Ali, S. Nageeb, David A. Miller, and David Yang, "Renegotiation-Proof Multilateral Enforcement," unpublished manuscript (2017).
- Andrews, Isaiah and Daniel Barron, "The Allocation of Future Business: Dynamic Relational Contracts with Multiple Agents," *American Economic Review* 106 no. 9 (2016): 2742–59.
- Baker, George., Robert Gibbons, and Kevin J. Murphy, "Subjective Performance Measures in Optimal Incentive Contracts," *Quarterly Journal of Economics* 109 no. 4 (1994): 1125–1156.
- Baker, George., Robert Gibbons, and Kevin J. Murphy, "Informal Authority in Organizations," *Journal of Law, Economics, and Organization* 15 no. 1 (1999): 56–73.
- Baker, George., Robert Gibbons, and Kevin J. Murphy, "Relational Contracts and the Theory of the Firm," *Quarterly Journal of Economics* 117 no. 1 (2002): 39–84.
- Baliga, Sandeep and Robert Evans, "Renegotiation in Repeated Games with Side-Payments," *Games and Economic Behavior* 33, no. 2 (2000): 159–176.
- Battigalli, Pierpaolo and Giovanni Maggi, "Rigidity, Discretion, and the Costs of Writing Contracts," *American Economic Review* 92, no. 4 (2002): 798–817.
- Battigalli, Pierpaolo and Giovanni Maggi, "Costly Contracting in a Long-Term Relationship," *The RAND Journal of Economics* 39, no. 2 (2008): 352–377.
- Barron, Daniel, Robert Gibbons, Ricard Gil, and Kevin J. Murphy, "Relational Adaptation under Reel Authority," *Management Science* 66 no. 5 (2019): 1868–89.
- Barron, Li, and Zator, "Morale and Debt Dynamics," unpublished manuscript (2019).
- Barron, and Andrews, "The Allocation of Future Business: Dynamic Relational Contracts with Multiple Agents," *American Economic Review* 106 no. 9 (2016): 2742–2759.
- Beshkar, Mostafa, "Arbitration and Renegotiation in Trade Agreements," *The Journal of Law, Economics, and Organization* 32, no. 3 (2016): 586–619.
- Board, Simon, "Relational Contracts and the Value of Loyalty," *American Economic Review* 101 no. 7 (2011): 3349–3367.
- Bull, Clive, "The Existence of Self-Enforcing Implicit Contracts," *Quarterly Journal of Economics* 102 no. 1 (1987): 147–159.

- Bull, Jesse and Joel Watson, "Evidence Disclosure and Verifiability," *Journal of Economic Theory* 118 (Issue 1, September 2004): 1–31.
- Bull, Jesse and Joel Watson, "Hard Evidence and Mechanism Design," *Games and Economic Behavior* 58 (2007): 75–93.
- Buzard, Kristy, "Self-Enforcing Trade Agreements and Lobbying," *Journal of International Economics* 108 (2017): 226–242.
- Chassang, Sylvain, "Building Routines: Learning, Cooperation and the Dynamics of Incomplete Relational Contracts," *American Economic Review* 100 no. 1 (2010): 448–465.
- Che, Yeon-Koo, and Seung-Weon Yoo, "Optimal Incentives for Teams," *American Economic Review* 91, no. 3 (2001): 525–541.
- den Haan, Wouter, Garey Ramey, and Joel Watson, "Contract-Theoretic Approaches to Wages and Displacement," *Federal Reserve Bank of St. Louis Review* 81(3) (1999).
- den Haan, Wouter, Garey Ramey, and Joel Watson, "Job Destruction and Propagation of Shocks," *The American Economic Review* 90 (2000): 482–498.
- Doornik, Kate, "Relational Contracting in Partnerships," *Journal of Economics and Management Strategy* 15 no. 2 (2006): 517–548.
- Dye, Ronald A., "Costly Contract Contingencies," *International Economic Review* 26 (1985): 233–250.
- Englmaier, Florian and Matthias Fahn, "Size Matters: How Over-Investments Relax Liquidity Constraints in Relational Contracts," *The Economic Journal* 129, no. 624 (2019): 3092–3106.
- Fahn, Matthias and Nicolas Klein, "Relational Contracts with Private Information on the Future Value of the Relationship: The Upside of Implicit Downsizing Costs," *American Economic Journal: Microeconomics* 11 no. 4 (2019): 33–58.
- Farrell, Joseph and Eric Maskin, "Renegotiation in Repeated Games," *Games and Economic Behavior* 1 no. 4 (1989): 327–360.
- Fong, Yuk-Fai and Jin Li, "Information Revelation in Relational Contracts," *The Review of Economic Studies* 84 no. 1 (2016): 277–299.
- Fuchs, William, "Contracting with Repeated Moral Hazard and Private Evaluations," *American Economic Review* 97 no. 4 (2007): 1432–1448.
- Fujiwara-Greve, Takako and Masahiro Okuno-Fujiwara, "Voluntarily Separable Repeated Prisoner's Dilemma," *The Review of Economic Studies* 76 (2009): 993–1021.
- Ghosh, Parikshit, and Debraj Ray, "Cooperation in Community Interaction Without Information Flows," *The Review of Economic Studies* 63 no. 3 (1996): 491–519.
- Gjertsen, Heidi, Theodore Groves, David A. Miller, Eduard Niesten, Dale Squires, and Joel Watson, "The Optimal Structure of Conservation Agreements and Monitoring," *Journal of Law, Economics, and Organization* (2021, in press).
- Goldlücke, Susanne and Sebastian Kranz: "Infinitely Repeated Games With Public Monitoring and Monetary Transfers," *Journal of Economic Theory* 147 no. 3 (2012): 1191–1221.

- Goldlücke, Susanne and Sebastian Kranz, “Renegotiation-Proof Relational Contracts,” *Games and Economic Behavior* 80 (2013): 157–178.
- Goldlücke, Susanne and Sebastian Kranz: “Reconciling Relational Contracting and Hold-up: A Model of Repeated Negotiations,” unpublished manuscript (2019).
- Grossman, Sanford J. and Oliver D. Hart, “The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration,” *Journal of Political Economy* 94 no. 4 (1986): 691–719.
- Grenadier, Steven R., Andrey Malenko, and Nadya Malenko, “Timing Decisions in Organizations: Communication and Authority in a Dynamic Environment,” *American Economic Review* 106 no. 9 (2016): 2552–81.
- Halac, Marina, “Relational Contracts and the Value of Relationships,” *American Economic Review* 102, no. 2 (2012): 750–79.
- Halonen, Maja, “Reputation and the Allocation of Ownership,” *Economic Journal* 112 issue 481 (2002): 539–58.
- Hart, Oliver D. and John Moore (1988), “Incomplete Contracts and Renegotiation,” *Econometrica* 56(4): 755–785.
- Holmström, Bengt, “Moral Hazard in Teams,” *The Bell Journal of Economics* (1982): 324–340.
- Iossa, Elisabetta and Spagnolo, Giancarlo, “Contracts as Threats: On a Rationale for Rewarding A While Hoping for B,” Brunel University, Department of Economics (2009).
- Itoh, Hideshi, and Hodaka Morita, “Formal Contracts, Relational Contracts, and the Threat-Point Effect,” *American Economic Journal: Microeconomics* 7, no. 3 (2015): 318–46.
- Klein, Daniel B. and Keith B. Leffler, “The Role of Market Forces in Assuring Contractual Performance,” *Journal of Political Economy* 89 no. 4 (1981): 615–641.
- Klimenko, Mikhail, Garey Ramey, and Joel Watson, “Recurrent Trade Agreements and the Value of External Enforcement,” *Journal of International Economics* 74 (2008, Issue 2): 475–499.
- Kostadinov, Rumen, “Renegotiation of Long-Term Contracts as Part of an Implicit Agreement,” unpublished manuscript (2019).
- Kvaløy, Ola and Trond E. Olsen, “Endogenous Verifiability and Relational Contracting,” *American Economic Review* 99, no. 5 (2009): 2193–2208.
- Legros, Patrick and Steven A. Matthews, “Efficient and Nearly-Efficient Partnerships,” *The Review of Economic Studies* 60, no. 3 (1993): 599–611.
- Levin, Jonathan, “Multilateral Contracting and the Employment Relationship,” *Quarterly Journal of Economics* 117 no. 3 (2002): 1075–1103.
- Levin, Jonathan (2003), “Relational Incentive Contracts,” *American Economic Review* 93(3): 835–857.
- Li, Jin, Niko Matouschek, and Michael Powell, “Power Dynamics in Organizations,” *American Economic Journal: Microeconomics* 9 no. 1 (2017): 217–241.
- Listokin, Yair, “Bayesian Contractual Interpretation,” *The Journal of Legal Studies* 39, no. 2 (2010): 359–374.

- Macaulay, Stewart, “Non-Contractual Relations in Business: A Preliminary Study,” *American Sociological Review* 28 (1963): 55–67.
- MacLeod, W. Bentley, “Optimal Contracting with Subjective Evaluation,” *American Economic Review* 93 no. 1 (2003): 216–240.
- MacLeod, W. Bentley, “Reputations, Relationships, and Contract Enforcement,” *Journal of Economic Literature* 4 (2007): 595–628.
- MacLeod, W. Bentley and Malcomson, James M., “Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment,” *Econometrica* 57 no. 2 (1989): 447–480.
- MacLeod, W. Bentley and Malcomson, James M. (1998), “Motivation and Markets,” *American Economic Review* 88(3): 388–411.
- Macneil, Ian R., “Contracts: Adjustment of Long-Term Economic Relations Under Classical, Neoclassical, and Relational Contract Law,” *Northwestern University Law Review* 72 no. 6 (1978): 854–904.
- Malcomson, James M., “Relational Incentive Contracts with Persistent Private Information,” *Econometrica* 84, no. 1 (2016): 317–346.
- Malcomson, James M., “Relational Incentive Contracts,” pp. 1014–1165 in *The Handbook of Organizational Economics*, edited by Robert Gibbons and John Roberts, Princeton, NJ: Princeton University Press (2013).
- Maskin, Eric and John Moore, “Implementation and Renegotiation,” *Review of Economic Studies* 66 (1999): 39–56.
- Miller, David and Joel Watson, “A Theory of Disagreement in Repeated Games with Bargaining,” *Econometrica* 81 (No. 6, November, 2013): 2303–2350.
- Pearce, David, “Renegotiation-Proof Equilibria: Collective Rationality and Intertemporal Cooperation,” Yale Cowles Foundation Discussion Paper 855 (1987).
- Pearce, David G. and Ennio Stacchetti, “The Interaction of Implicit and Explicit Contracts in Repeated Agency,” *Games and Economic Behavior* 23 no. 1 (1998): 75–96.
- Radner, Roy, “Repeated Principal-Agent Games with Discounting,” *Econometrica* 53 no. 3 (1985): 1173–1198.
- Ramey, Garey and Joel Watson, “Contractual Fragility, Job Destruction, and Business Cycles,” *Quarterly Journal of Economics* 112 (1997): 873–911.
- Ramey, Garey and Joel Watson, “Bilateral Trade and Opportunism in a Matching Market,” *Contributions to Theoretical Economics* 1 no. 1 (2001), Article 3.
- Ramey, Garey and Joel Watson, “Contractual Intermediaries,” *Journal of Law, Economics, and Organization* 18 (2002): 362–384.
- Rauch, James and Joel Watson, “Starting Small in an Unfamiliar Environment,” *International Journal of Industrial Organization* 21 (2003): 1021–1042.
- Ray, Debraj, “The Time Structure of Self-Enforcing Agreements,” *Econometrica* 70 no. 2 (2002): 547–582.
- Rayo, Luis, “Relational Incentives and Moral Hazard in Teams,” *Review of Economic Studies* 74 no. 3 (2007): 937–963.

- Rubinstein, Ariel, "An Optimal Policy for Offenses that May Have Been Committed By Accident," in *Applied Game Theory*, ed. by S. Brams, A. Schotter and G. Schwodiauer, Wurzburg: Physica-Verlag (1979): 406–413.
- Rubinstein, Ariel, "Strong Perfect Equilibrium in Supergames," *International Journal of Game Theory* 9, no. 1 (1980): 1–12.
- Schmidt, Klaus M. and Monika Schnitzer, "The Interaction of Explicit and Implicit Contracts," *Economics Letters* 48 no. 2 (1995):193–199.
- Schöttner, Anja, "Relational Contracts, Multitasking, and Job Design," *Journal of Law, Economics, and Organization* 24 no. 1 (2008): 138–162.
- Schwartz, Alan and Joel Watson, "The Law and Economics of Costly Contracting," *Journal of Law, Economics, and Organization* 20 (April 2004): 2–31.
- Schwartz, Alan and Joel Watson, "Conceptualizing Contractual Interpretation," *Journal of Legal Studies* 42 (2013, Issue 1): 1–34.
- Shapiro, Carl and Joseph Stiglitz, "Equilibrium Unemployment as a Worker Discipline Device," *American Economic Review* 74 no. 3 (1984): 433–444.
- Shavell, Steven, "On the Writing and the Interpretation of Contracts," *The Journal of Law, Economics, and Organization* 22, no. 2 (2006): 289–314.
- Sobel, Joel, "For Better or Forever: Formal Versus Informal Enforcement," *Journal of Labor Economics* 24 no. 2 (2006): 271–297.
- Spear, Stephen E., and Sanjay Srivastava, "On Repeated Moral Hazard with Discounting," *The Review of Economic Studies* 54, no. 4 (1987): 599–617.
- Telser, Lester G., "A Theory of Self-Enforcing Agreements," *Journal of Business* 53 no. 1 (1980): 27–44.
- Thomas, Jonathan and Tim Worrall, "Dynamic Relational Contracts Under Complete Information," *Journal of Economic Theory* 175 (2018): 624–651.
- Watson, Joel, "Starting Small and Renegotiation," *Journal of Economic Theory* 85 (1999): 52–90.
- Watson, Joel, *Strategy: An Introduction to Game Theory*, New York, NY: W.W. Norton and Company (2002).
- Watson, Joel, "Starting Small and Commitment," *Games and Economic Behavior* 38 (2002): 176–199.
- Watson, Joel, "Contract, Mechanism Design, and Technological Detail," *Econometrica* 75 no. 1 (2007): 55–81.
- Watson, Joel, "Contract and Game Theory: Basic Concepts for Settings with Finite Horizons," *Games* 4 (2013): 457–496.
- Watson, Joel, David Miller, and Trond Olsen, "Relational Contracting, Negotiation, and External Enforcement," *The American Economic Review* 110 no. 7 (2020): 2153–2197.
- Williamson, Oliver, *The Economic Institutions of Capitalism*, New York, NY: The Free Press (1985).
- Yang, Huanxing, "Nonstationary Relational Contracts with Adverse Selection," *International Economic Review* 54 no. 2 (2013): 525–547.