Relational Contracts: Methodological Overview

Joel Watson, Spring 2024

Contents

1 Contract Basics

1.1 The definition of contract and modes of enforcement

Fundamental notion/definition (broad view)...

A *contract* is an agreement that is intended to be enforced.

(Intent of the contracting parties, external parties, or society...)

Examples include:

- Employment and independent-contractor agreements between workers and firms,
- Collaboration and supply-chain agreements between firms,
- Procurement agreements between government agencies and firms,
- Marriage and prenuptial agreements between domestic partners,
- Rental agreements between landlords and tenants,
- Financial agreements such as loans between creditors and debtors,
- International agreements between countries (environmental, trade, security),
- Community-level agreements on production, lending, and support, etc.

These are all contracts to the extent that the parties expect and intend their agreements to be enforced.

We seek to model:

- *Contract formation* and *renegotiation*, typically through negotiation;
- *Productive interaction*, which involves incentive problems (moral hazard) and welfare implications;
- *Information*, typically regarding problems of observability and verifiability; and
- *Enforcement*.

Modes of 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.

Almost all contracts are enforced by a combination of self-enforcement and external enforcement.

Every contract can be described as combining:

- An *external part (external contract)*, which refers to contractual provisions that instruct the external enforcer on how to intervene in the relationship, typically by compelling monetary transfers (such as penalties imposed by the court) as a function of information available to the enforcer; and
- An *internal part (internal contract)*, recording how the contracting parties have agreed to themselves act.

Unhelpful terminology in the literature:

- Written versus oral/non-written,
- Implicit versus explicit,
- Formal versus informal, and
- Legal versus nonlegal.

We seek to understand for any given contractual setting:

- Whether efficient outcomes are reached or can be reached in principle,
- The nature of and implications for externalities,
- The resulting allocation of resources,
- The form of contracts and exercise of enforcement,
- Implications of the structure and constraints of enforcement technologies,
- Implications of informational imperfections,
- The relation between modes of enforcement, and
- Other economic considerations important for specific applications.

1.2 Relational contracts

- Researchers tend to use the term “relational contract” in reference to contractual relationships that rely on self-enforcement and persist over time.
- “Relational incentive contract” is the subset of relational contracts that entail incentive problems (moral hazard). This is the setting we will focus on.
- The most intensively studied class of relational contracts are employment relationships.

1.3 Game-theoretic foundations

Contractual settings are strategic settings, and therefore the analysis of contract requires game-theoretic reasoning.

- I 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.
- Events that all of the players see are called *observable*, and information available to an external enforcer is called *verifiable*.

Key foundations for analysis of relational contracts:

- Repeated game theory,
- Bargaining theory,
- Information theory, and
- Mechanism design.

1.4 Historical notes

"Relational contract" in the legal literature:

- Scholars observed 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.

Progression toward theories of relational contract in the labor literature:

- Baily (1974) and Azariadis (1975) recognize that an employment relationship rather than spot market exchange can manage risk sharing, and call it an "implicit contract." Rosen (1985) reviews the literature. Key aspects of contracting and enforcement are not modelled.

- Thomas and Worrall (1988) model self-enforcement of the relational contract in this context. Logic:

  - Worker-firm matched pair can produce $y \sim \sigma$ in a given discrete period.
  - Spot labor market: firm gets 0 (free entry), worker gets $u_1(y)$.
  - In a worker-firm long-term relationship, parties unilaterally decide whether to separate and go to the spot market forever after.
  - Relationship is stable and efficiently shares risk with a constant wage m_1 if $\frac{1}{1-\delta}u_1(m_1) \geq \frac{\delta}{1-\delta}E_{z\sim\sigma}u(z)$ and $y - m_1 + \frac{\delta}{1-\delta}(\overline{y} - m_1)$ for all $y \in \operatorname{supp}\sigma$.

Early work on incentives and relationships in IO and labor:

- 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.

2 Emerging Theoretical Foundations of RC

(not comprehensive)

Early work in game theory:

- Mechanics of self-enforcement in relational contracts given by the conditions for equilibria in repeated games; core elements 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.

Key papers (for modeling elements in particular) since the early 1980s:

- Radner (1985) developed an infinite-horizon repeated-game model of a principal-agent setting with external enforcement of short-term contracts (transfer contingent on output, assumed observable and verifiable). He studied equilibria in "review strategies," showing that near-efficient equilibria exist with low discounting. Stages in each discrete period:
  - Principal sets a bonus scheme (b_1^1, b_1^0), with $b_1^1 \geq b_1^0 \geq 0$.
  - Agent chooses effort $e \in [0, 1]$ privately (principal does not observe, not verifiable).
  - Revenue is $r = 1$ with probability e and $r = 0$ with probability $1 - e$.
  - Revenue is observed and verifiable. Bonus contract externally enforced.
  - Agent's payoff is $u_1(b_1^r) - c(e)$; principal's payoff is $r - b_1^r$.

- Spear and Srivastava (1987) examined equilibria more generally and put the analysis in terms of a dynamic program featuring continuation values.
- We now see that similar conclusions hold without external enforcement.

- MacLeod and Malcomson (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 deeper analysis of contractual arrangements than Shapiro and Stiglitz (1984) studied, including transfer incentives, and an observation about pooled effort constraints. Quasilinear utility simplifies the analysis. Efficiency is achieved for low discounting. Stages in each discrete period:
  - Simultaneously, worker and firm choose whether to unilaterally separate.
  - If together, worker chooses effort $e \in [0, 1]$ and firm observes e (not verifiable).
  - Firm chooses any bonus $b_1 \geq \underline{b}_1$, where $\underline{b}_1$ is the externally enforced wage.
  - Worker's payoff is $b_1 - c(e)$; firm's payoff is $y(e) - b_1$.

- Levin (2001) expanded MacLeod and Malcomson's (1989) model to study a range of production technologies, to further characterize equilibrium strategies, and ostensibly add bargaining power. Results are similar; separation is shown to not be needed for punishment. Stages in each discrete period:
  - Principal offers a spot contract consisting of externally enforced transfer $m = (m_1, m_2)$ and (cheap talk) schedule of voluntary bonuses b_1 contingent on observed output.
  - Agent accepts or rejects the offer.
  - If agent accepted, then the players are engaged for the period, agent observes a random draw and selects effort e, stage payoff vector u depends on effort and move of nature, and an unverifiable public signal is realized. If agent rejected, then the players are disengaged for the period and receive outside-option payoff vector $\underline{u}$.
  - If engaged, principal chooses bonus b_1.
  - The payoff vector in the period is $m + u + (b_1, -b_1)$ if engaged and $\underline{u}$ if disengaged.

- 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. BGM provide logic to suggest that increasing the effectiveness of external enforcement undermines cooperation. Stages in each discrete period:
  - Firm offers a spot contract consisting of a wage m_1, bonus b_1^{vp} contingent on a verifiable performance measure, and discretionary bonus b_1^{d} promised contingent on an observed but unverifiable signal.
  - Worker accepts or rejects the offer.
  - If worker accepted, then the players are engaged for the period, agent selects two-dimensional effort e, performance and signal are realized, and the players receive stage payoff vector $u(e)$. If agent rejected, the players are disengaged for the period and receive outside-option payoff vector $\underline{u}$.
  - If engaged, firm observes the performance measure and signal, and chooses bonus b_1^{d}.
  - Payoff vector is $m + u + (b_1^{\mathrm{vp}}, -b_1^{\mathrm{vp}}) + (b_1^{\mathrm{d}}, -b_1^{\mathrm{d}})$ if engaged and $\underline{u}$ if disengaged.

- Che and Yoo (2001) examined relational contracting in the context of team production, with a manager setting a stationary external contract for the team. They find that joint performance evaluation may dominate relative performance evaluation, because it creates better incentives for the workers to punish each other following deviations. Kvaløy and Olsen (2006) studied team production with unverifiable effort signals, so that payments from the manager must be self-enforced, and they considered collusion by the workers.

- Ramey and Watson (1997, 2001) and den Haan, Ramey, and Watson (2000) put relational contracts in the context of a frictional matching market, showing how incentives in employment relationship interact across markets in the presence of shocks.

- Ghosh and Ray (1996) look at frictions that arise due to private information in the matching market. Kandori (1992), Kranton (1996), Fujiwara-Greve and Okuno-Fijiwara (2009), Deb (2020), and others put relational contracts in the context of random matching in a community. These papers suggest some sort of community norm for the equilibrium in the overall game.

- The full-blown recursive characterization of "perfect public equilibrium" values in general repeated games was pioneered by Abreu, Pearce, and Stacchetti (1990).

- Goldlücke and Kranz (2012, 2013) characterized perfect public equilibrium values for general repeated games with quiasilinear utility and separate phases for transfers and productive actions in each period, which is a general platform for relational-contracting models.

- Miller and Watson (2013) examine relational contracting with an explicit account of bargaining and proposed the notion of "contractual equilibrium" in general settings with self-enforcement.

- Watson, Miller, and Olsen (2020) extended the theory to settings with external enforcement of long-term contracts.

Additional points:

- The above summary does not include the significant amount of research on settings with incomplete information, such as where parties have private information about costs, benefits, and outside options. Examples: Watson (1999), Halac (2012), Malcomson (2016), Li and Matouschek (2013), Hua and Watson (2022). We'll look briefly at this area later in the week if there is interest.

- There are good opportunities to contribute to the theory literature, such as to build more general models of relational contracting with outside options, matching, short-run and persistent shocks; to investigate alternative theories of negotiation; to investigate the interaction between contracting at various level (bilateral relationships, the community, and such); and to examine multilateral contracting.

- There are good opportunities to contribute to the applied literature, such as to build a macro model of the labor market with an evolving distribution of characteristics (worker, firm, or match type), with or without private information, or to study optimal nonstationary contracts in team production.

3 Production Stage – Notation

Consider a general class of settings with the following properties: quasi-linear utility (that is, 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 described by an *underlying game* $\gamma = (A, X, \lambda, u, P)$, representing the players' productive actions, their personal costs and benefits, and the intrinsic distribution of returns, 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 $\lambda\colon A \to \Delta X$,
- a payoff function $u\colon A \to \mathbb{R}^n$, 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.

The realized outcome $x \in X$ is commonly observed by the players, but only the partition element that contains x, denoted by $P(x)$, 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, and extend u to the space of mixed actions. Player i observes only x and her own action a_i.[1]

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

The external part of the players' contract is a transfer function $b \colon X \to \mathbb{R}_0^n$.

$\mathbb{R}_0^n = \{m \in \mathbb{R}^n \mid \sum_{i=1}^n m_i = 0\}$ is the space of *balanced transfers* (components sum to zero).

For any 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

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

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

Function b is constrained to be P-measurable, because the external enforcer can observe only what is verifiable about the outcome. Thus, $x \in P(x')$ imposes the requirement that $b(x) = b(x')$.

4 Examples of Productive Interaction

4.1 Principal-agent example

Consider a relationship between an agent (worker, player 1) and a principal (manager, player 2). In the underlying game, the worker chooses effort $a_1 \geq 0$, the outcome is $x = a_1$, and the payoff vector is $u(a_1) = (-a_1^2, a_1 + a_1^2)$.

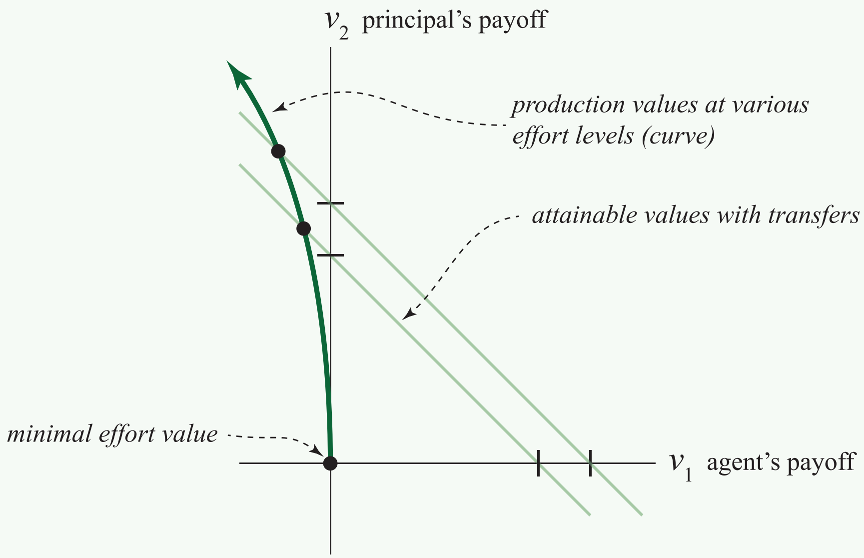

v_2 principal's payoff
production values at various effort levels (curve)
attainable values with transfers
minimal effort value
v_1 agent's payoff

If the outcome is verifiable, then any effort level $\hat{a}_1$ can be achieved with a contract that specifies $b(a_1) = (0, 0)$ for $a_1 < \hat{a}_1$ and $b(a_1) = (\hat{a}_1^2 + p, -\hat{a}_1^2 - p)$ for $a_1 \geq \hat{a}_1$, for any $p \in [0, \hat{a}_1]$.

If the outcome is not verifiable, then external enforcement is useless and effort must be 0.

Turn to relational contracting…

4.2 Project-choice example

This is another principal-agent setting, except the worker (player 1) chooses not just whether to exert high effort but also to which of three projects to apply his effort. He can expend effort on only one project.

The manager (player 2) observes the worker's effort choice and receives the revenue that it generates. The manager has no action in the underlying game.

The outcome includes a noisy binary signal of the worker's effort.

The set of feasible effort choices is $A_1 = \{0, 1, 2, 3\}$, where $a_1 = 0$ represents no effort, $a_1 = 1$ means applying effort to project 1, $a_1 = 2$ means applying effort to project 2, and $a_1 = 3$ means applying effort to project 3.

The signal is 1 with probability $\sigma(a_1)$ and 0 with probability $1 - \sigma(a_1)$.

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

Note that $a_1 = 1$ is the efficient effort choice, yielding a joint value of 8.

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.[2]

The signal is verifiable but player 1's effort choice is not verifiable, so the outcome partition is

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

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.

[2] Note that the contingent distribution function λ is given by $\lambda(0)(00) = 1$, $\lambda(1)(10) = 1/2$, $\lambda(1)(11) = 1/2$, $\lambda(2)(20) = 3/4$, $\lambda(2)(21) = 1/4$, $\lambda(3)(30) = 0$, and $\lambda(3)(31) = 1$.

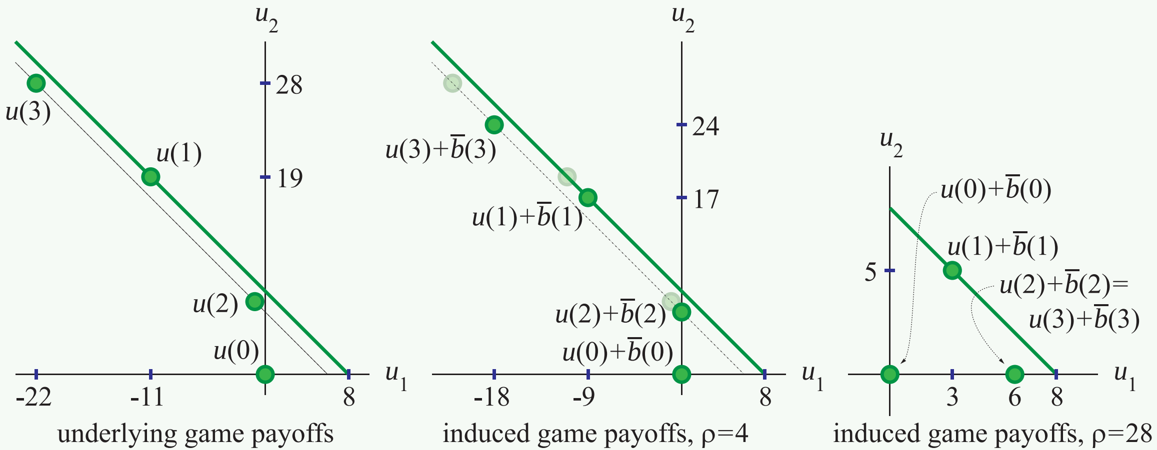

Figure 1: Payoffs in the project-choice example.

Because only player 1 has an action in the stage game, it is easy to visualize player 1's incentives by looking at a 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 $u(a_1) + \bar{b}(a_1)$ the right-most point in the graph.

Four things to note:

- 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.

- The players cannot do better because no contract can implement the efficient action $a_1 = 1$. This is easy to see by noting that $u_1(0) + \bar{b}(0) > u_1(1) + \bar{b}(1)$ for $\rho < 22$, $u_1(2) + \bar{b}(2) > u_1(1) + \bar{b}(1)$ for $\rho < 40$, and $u_1(3) + \bar{b}(3) > u_1(1) + \bar{b}(1)$ for $\rho > 22$.

- The difference between $\max\{u_1(a_1) + \bar{b}(a_1) \mid a_1 = 0, 2, 3\}$ and $u_1(1) + \bar{b}(1)$, though strictly positive, is minimized by choosing $\rho = 28$ as shown in the right graph of Figure 1.

- If the players' contract specifies $\rho = 4$, then the induced game has two Nash equilibria, $a_1 = 0$ and $a_1 = 2$.

Items three and four will play an important role in the analysis of relational contracting later.

4.3 Team-production example

The underlying game models interaction between two coworkers (players 1 and 2) who simultaneously contribute effort to a joint project. $A_1 = A_2 = \{0, 1\}$. A player who chooses high effort pays a personal cost of 4.

In the end, the joint project is either a success, which pays both players 10, or a failure, which pays them nothing. Success occurs with probability 0.8 if $a_1 = a_2 = 1$ (both players chose high effort), with probability 0.5 if $a_1 + a_2 = 1$ (exactly one of the players chose high effort), and with probability zero if $a_1 = a_2 = 0$.

The players do not observe each others' effort choices but they commonly observe whether the project is a success.

In the case of a failed project, a commonly observed noisy signal is generated about who is responsible. The signal identifies player 1 with probability $\frac{1}{2} - \frac{1}{2}\sigma(a_1 - a_2)$ and player 2 with probability $\frac{1}{2} - \frac{1}{2}\sigma(a_2 - a_1)$, where $\sigma \in [0, 1]$ is a parameter that measures the accuracy of the signal.

Note that if $\sigma = 0$ then the signal provides no information. If $\sigma = 1$ then it perfectly distinguishes between $a = (0, 1)$ and $a = (1, 0)$ but is still noisy in differentiating these action profiles from $a = (1, 1)$.

The outcome space is thus $X = \{1, 01, 02\}$, where $x = 1$ denotes a successful project, $x = 01$ denotes a failed project with player 1 identified as responsible, and $x = 02$ denotes failure with player 2 identified as responsible.

Recall that we write the underlying-game payoffs as a function of the action profile, averaging over the outcomes.

$$
\begin{aligned}
u(1,1) &= 0.8 \cdot (10,10) - (4,4) = (4,4)\\
u(0,1) &= 0.5 \cdot (10,10) - (0,4) = (5,1)\\
u(1,0) &= 0.5 \cdot (10,10) - (4,0) = (1,5)\\
u(0,0) &= 0 \cdot (10,10) - (0,0) \quad = (0,0).
\end{aligned}
$$

With no verifiability, where P is the coarsest partition, external enforcement is useless but the players can achieve some measure of success. They coordinate on one of the asymmetric equilibria of the underlying game, $(0, 1)$ or $(1, 0)$, which yields a joint value of 6 that they divide via an up-front transfer.

With full verifiability, where P is the finest partition, the players may fare better depending on the parameter σ.

In the case of $\sigma > 0$, externally enforced transfers can align incentives by tailoring the transfer to the signal of responsibility.

Suppose that the external contract specifies $b(1) = (0, 0)$, $b(01) = (-\rho, \rho)$, and $b(02) = (\rho, -\rho)$, where $\rho > 0$ is a penalty that the party held responsible must pay to the other player. Then the induced-game expected payoffs are

$$\begin{aligned}
u(1,1) + \bar{b}(1,1) &= (4,4) \\
u(0,1) + \bar{b}(0,1) &= (5,1) + (0.5)\sigma(-\rho,\rho) \\
u(1,0) + \bar{b}(1,0) &= (1,5) + (0.5)\sigma(\rho,-\rho) \\
u(0,0) + \bar{b}(0,0) &= (0,0) \qquad .
\end{aligned}$$

With ρ set sufficiently high, action profile $(1, 1)$ becomes an equilibrium of the induced game.

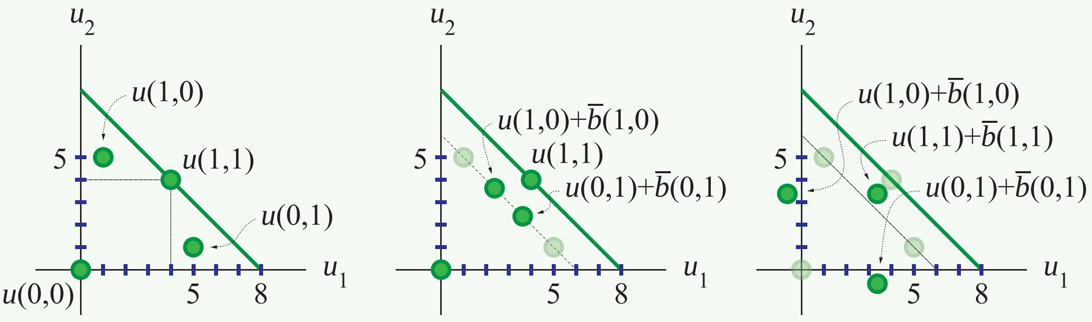

underlying game payoffs

induced game payoffs, full verifiability, $\sigma > 0$

induced game payoffs, full verifiability, $\sigma = 0$, unbalanced transfers

In the case of $\sigma = 0$, externally enforced transfers do nothing for incentives because the induced game would be exactly the same as the underlying game.

This is a case where the players could benefit from a commitment to extract money from them jointly. For instance, suppose that a contract specifying $b(1) = (0, 0)$ and $b(01) = b(02) = (-\rho, -\rho)$ for $\rho > 0$ is externally enforced. Then the induced-game expected payoffs are:

$$\begin{aligned} u(1,1) + \bar{b}(1,1) &= (4,4) + (0.2)(-\rho,-\rho) \\ u(0,1) + \bar{b}(0,1) &= (5,1) + (0.5)(-\rho,-\rho) \\ u(1,0) + \bar{b}(1,0) &= (1,5) + (0.5)(-\rho,-\rho) \\ u(0,0) + \bar{b}(0,0) &= (0,0) + (-\rho,-\rho) \qquad . \end{aligned}$$

By setting $\rho \geq 10/3$, action profile $(1, 1)$ becomes an equilibrium.

This is Holmström's (1982) point that if output can be verified but not individuals' efforts, achieving efficient effort incentives requires team punishment when output falls below the target level.

4.4 General results for team production with verifiable revenue

Consider a version of the team-production model in which revenue is fully verifiable and nothing more, which is the focus of Holmström (1982) and the more general modeling of Legros and Matthews (1993).

Underlying game: n team members simultaneously exert effort toward a joint project. Player i chooses effort level $a_i \geq 0$ at personal cost $c_i(a_i)$, measured in monetary terms. The outcome is defined as the effort profile, so that $x = a$ with probability one, and $A = X = \mathbb{R}^n_+$.

The project's revenue is given by a function $f\colon \mathbb{R}^n_+ \to \mathbb{R}_+$.

Suppose that player 1 receives the revenue, which means $u_1(a) = f(a) - c_1(a_1)$ and $u_i(a) = -c_i(a_i)$ for $i \neq 1$.

Partition P represents that the external enforcer can verify x and nothing more. Thus,

$$P = \{\{a \mid f(a) = x\} \mid x \in \mathbb{R}_+\}\,.$$

An external contract therefore specifies a transfer between the players as a function of revenue x.

Assume all functions are differentiable; $c_i', c_i'' > 0$ and $\lim_{a_i \to \infty} c_i'(a_i) = \infty$; and f is strictly increasing in all arguments and concave.

Let a^* be the unique efficient effort profile, solving

$$\max_a f(a) - \sum_{i=1}^{n} c_i(a_i).$$

Theorem: In the team-production setting under the assumptions stated above and with external enforcement of only balanced transfers (meaning $b\colon A \to \mathbb{R}_0^n$), effort profile a^* is not implementable and the contractual equilibrium is inefficient. However, if the players could commit to imbalanced transfers, allowing $b(a) < 0$, then a^* is implementable and the contractual equilibrium is efficient.

Legros and Matthews (1993) examine a more general setting. The outcome is $x = a$ as before, but now payoffs are any functions of the effort profile and partition P is unrestricted (not defined in relation to revenue).

$$Q^i(\hat{a}) \equiv \bigcup_{a_i \in A_i} P(a_i, \hat{a}_{-i})$$

and

$$w_i(a, \hat{a}) \equiv \begin{cases} \max\{u_i(a_i', \hat{a}_{-i}) \mid (a_i', \hat{a}_{-i}) \in P(a)\} - u_i(\hat{a}) & \text{if } Q^i(\hat{a}) \cap P(a) \neq \emptyset \\ -\infty & \text{if } Q^i(\hat{a}) \cap P(a) = \emptyset \end{cases}.$$

Theorem: Effort profile $\hat{a}$ is implementable if and only if $\sum_{i=1}^n w_i(a, \hat{a}) \leq 0$ for every $a \in A$.

If an efficient effort vector cannot be achieved with external enforcement, turn to relational contracting…

Start with a review of recursive method for analysis of repeated games…

5 Recursive Method for Analysis of Repeated Games

Let us consider an n-player, infinitely repeated game, with stage game $\langle A, X, \lambda, u, P \rangle$ and a common discount factor $\delta \in (0, 1)$.

The elements of the stage game are as defined before.

Normalize payoffs in each period by multiplying by $1 - \delta$.

5.1 Definitions and assumptions

- Set of public histories: $H \equiv \cup_{k=0}^{\infty} X^k$, where X^0 is the null history at the start of the game. Set of personal histories for player i: $H^i \equiv \cup_{k=0}^{\infty} (A_i \times X)^k$.

- A pure strategy for player i is a function $s_i \colon H^i \to A_i$ and a mixed (behavior) strategy is $\sigma_i \colon H^i \to \Delta A_i$. If the strategy depends only on the public history, then write $s_i \colon H \to A_i$ and $\sigma_i \colon H \to \Delta A_i$. Strategy profiles: s and σ.

- We usually assume players have access to a public randomization device that delivers a random draw at the end of each period and on which they can condition their future behavior. It can formally be incorporated into the specification of X and λ, or included separately.

- Given a strategy σ_i and a history $h \in H^i$, the *continuation strategy* from h, denoted $\sigma_i|h$, is defined by $\sigma_i|h(h') = \sigma_i(hh')$ for all $h' \in H^i$, where "hh'" denotes h concatenated with the sequence h'.

- For the equilibrium definition below, we will constrain attention to strategy profiles in which the players condition on only the public history.

- Given σ and any t-period public history $h \in H$, the *continuation value* from h is given by

$$v(\sigma|h) \equiv \mathrm{E}_\sigma \left[\sum_{k=1}^{\infty} \delta^{k-1}(1-\delta)u(a^{t+k}) \right],$$

  where $\{a^{t+k}\}_{k=1}^{\infty}$ is the sequence of action profiles from period $t+1$ on, and the expectation is taken with respect to the strategy profile σ.

- A strategy profile σ^* is called a *perfect public equilibrium* (PPE, an extension of subgame perfection) if (i) it conditions on only the public history and (ii) for every $h \in H$, every player i, and every alternative strategy σ_i', it is the case that $v_i(\sigma^*|h) \geq v_i(\sigma_i', \sigma_{-i}^*|h)$.

- Single deviation property: In relation to a strategy σ_i and history $h \in H$, call σ_i' a *single deviation at* h if $\sigma_i'(h') = \sigma_i(h')$ for all $h' \neq h$.

Theorem: A strategy profile σ^* is a perfect public equilibrium if and only if it conditions on only the public history and no player can gain by a single deviation after any history.

5.2 Recursive structure of continuation values

For any history $h \in H$ and strategy σ, let $\alpha \equiv \sigma(h) \in \Delta A$ be the mixed action profile that σ would play in the next period.

Note that

$$v(\sigma|h) = (1-\delta)u(\alpha) + \delta \mathrm{E}_{\lambda(\alpha)}\left[v(\sigma|hx)\right].$$

For any set of continuation values W from the start of the next period, let $y\colon X \to W$ describe the continuation value that the players coordinate on as a function of the outcome x in the current period.

Then define $\overline{y}\colon A \to \mathbb{R}^2$ to be the expected value as a function of the action profile played in the current period: $\overline{y}(a) \equiv E_x[y(x)\,|\,x \sim \lambda(a)]$.

Incorporating the anticipated continuation value, interaction in the stage game of the current period is effectively the induced static game

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

where A is the set of action profiles, and payoffs are the convex combination of stage-game payoffs and continuation values.

The players can self-enforce any mixed action profile $\alpha \in \Delta A$ that is a Nash equilibrium of this induced game, resulting in the following continuation value from the action phase in the current period.

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

Given $\langle A, X, \lambda, u, P\rangle$ and $y\colon X \to \mathbb{R}^n$, call action profile $\alpha \in \Delta A$ **enforced relative to $\boldsymbol{y}$** if it is a Nash equilibrium of the induced game.

Call w **supported relative to $\boldsymbol{W}$** if there is a function $y\colon X \to W$ and an action profile $\alpha \in \Delta A$ that is enforced relative to y, such that Equation 1 holds.

For any $W \subset \mathbb{R}^n$, define

$$B(W) \equiv \{w \in \mathbb{R}^n \mid w \text{ is supported relative to } W\}.$$

Let V denote the set of perfect public equilibrium payoff vectors.

Call any set $W \subset \mathbb{R}^n$ **self-generating** if $W \subset B(W)$.

Theorem: Assume that A is finite. Then the following are true: (a) B is monotone. (b) If W is self-generating then $W \subset V$. (c) V is the largest self-generating set.

Theorem: Assuming A is finite, the following are true:

- If W is compact then $B(W)$ is compact.
- $W \subset B(W)$ implies $W \subset \cap_{k=1}^{\infty} B^k(W) \subset V$.
- If W is bounded and $V \subset W$ then $\cap_{k=1}^{\infty} B^k(W) = V$.
- V is compact. Also, if λ incorporates a suitable public-randomization device then, for all $v \in V$ there is a perfect public equilibrium whose continuation values are extreme points of V.
- For $\delta_1 < \delta_2$, $V(\delta_1) \subset V(\delta_2)$, where these are the sets of equilibrium payoff vectors in the case of δ_1 and δ_2.

6 Relational Incentive Contract

Components of a general framework:

- Discrete-time interaction with an infinite number of periods, $t = 1, 2, 3, \ldots$;
- Stationary productive environment, underlying stage game (A, X, λ, u, P);
- External enforcement of monetary transfers given by a P-measurable function $b \colon X \to \mathbb{R}^n_0$ that the external contract specifies (in general, could depend on history);
- Multiple-phase interaction within a period:
  - The *negotiation phase*, where players form their contract and make a transfer $m \in \mathbb{R}^n_- \equiv \{m \in \mathbb{R}^n \mid m_1 + m_2 + \cdots + m_n \leq 0\}$, followed by
  - The *production phase*, where the players choose actions in the underlying game; and
- Expected payoffs in period t given by $m^t + u(a^t) + \overline{b}^t(a^t)$, normalized by $1 - \delta$, where δ is the common discount factor.

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.

Applications include principal-agent relationships (worker-manager, subcontractor-firm, etc.), production by teams (work groups within a firm, partnerships, joint-ventures, etc.), relationships between regulators and firms, international agreements, and conservation agreements.

These are all relationships in which, in each period of time, the contracting parties individually choose productive actions (such as effort levels) and can also make monetary transfers. They form contracts to manage their incentives.

Some 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 severe 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 may comment on these variations later.

Alternatives for modeling negotiation:

- Choice of cooperative or noncooperative bargaining models.
- For noncooperative models, variety of alternative assumptions regarding how communication in the negotiation phase translates into coordinated play in the production phase of the current period and beyond.

Leave the negotiation phase unspecified for now, but note that the payoff relevant aspect of play in the negotiation phase is an immediate transfer that the players make in reaching an agreement.

7 Recursive Method Review: Incentives in Production Phase

Let’s explore the players’ incentives in the production phase of any given period, which we refer to as “the current period” and drop the t superscripts.

We’ll adapt the method pioneered by Abreu, Pearce, and Stacchetti (1990).

Let y denote the continuation value that the players coordinate to obtain from the start of the next period, as a function of the current-period outcome x, and let $\overline{y}$ give the expectation as a function of the current-period action profile a.

Interaction in the production phase of the current period is essentially play of the induced game

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

Self-enforcement in the current period amounts to coordination on a Nash equilibrium of this induced static game.

Similar to the analysis of standard repeated games, 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 from the current-period production phase,

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

For any P-measurable function $b\colon X \to \mathbb{R}^n_0$ and any set Y of functions from A to $\mathbb{R}^n$, define

$$\begin{aligned} D(b,Y) \equiv \Big\{ (1-\delta)\left[u(\alpha)+\bar{b}(\alpha)\right]+\delta\overline{y}(\alpha) \;\Big|\; & y \in Y \text{ and} \\ & \alpha \text{ is a Nash equilibrium of } \langle A, (1-\delta)(u+\bar{b})+\delta\overline{y}\rangle \Big\}. \end{aligned}$$

This operator is a generalization of the APS operator.

8 Settings with Trivial External Enforcement, Inactive Contracting

Consider settings in which the externally enforced transfer function b is exogenously fixed and the same in every period (no external enforcement, or the enforcement technology requires commitment to a single transfer function over time). Without loss of generality, assume that $b \equiv \mathbf{0}$.

Two alternative models of the negotiation phase:

- Active contracting and
- Inactive contracting.

Consider the latter.

In the this version of the model, all that happens in the negotiation phase is the players simultaneously make voluntary transfers (publicly observed), modeled noncooperatively. The vector sum of transfers defines $m \in \mathbb{R}^n_-$, the total transfer in the negotiation phase. There is no real negotiation accounted for in the negotiation phase.

Quite a few relational-contract models essentially fall into this category, including MacLeod and Malcomson (1989) and Levin (2003) on principal-agent relationships, and Doornik (2006) and Schöttner (2008) on partnerships/team production.

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, and let W' be the set of PPE continuation values from the production phase in each period.

Let $F(W) \equiv \{y\colon X \to \operatorname{conv} W\}$. Note that

$$W' = D(\mathbf{0}, F(W^*))$$

and

$$W^* = \operatorname{tri} W',$$

where

$$\operatorname{tri} W \equiv \Big\{(1-\delta)m + w \;\Big|\; m \in \mathbb{R}^n_-, w \in W, \text{ and for every } i, \\ \text{there exists } \underline{w}^i \in W \text{ such that } (1-\delta)m_i + w_i \geq \underline{w}^i_i\Big\}.$$

Because operators D and tri are monotone, the composition $\mathrm{tri}\, D(\mathbf{0}, F(\cdot))$ is also monotone and so we have the following result:

Theorem: The set of PPE payoff vectors in the relational-contract game is exactly the largest fixed point of the operator $\mathrm{tri}\, D(\mathbf{0}, F(\cdot))$, denoted W^*.

Here is where transferable utility, the simplifying assumption of relational-contract theory made by MacLeod and Malcomson (1989) and many others since, starts to deliver benefits in the characterization of equilibrium.

Observe that, whatever is W, $\operatorname{tri} W$ is a generalized triangle with a linear frontier of slope -1.

If W is compact, there are numbers $\underline{w}_1^1, \underline{w}_2^2, \ldots, \underline{w}_n^n$ such that, letting

$$L = \max_{w \in W} \sum_{i=1}^{n} w_i$$

denote the *level*, we have $v \in \operatorname{tri} W$ 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 $\underline{w}_i^i$.

A vector is in W^* if and only if is can be expressed as the players jointly getting level L^* and making a monetary transfer that is unrestricted except for each player receiving at least her minimum 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; for relatively simple production technologies, it becomes straightforward to calculate. Furthermore, we do so without having to describe the equilibrium strategies.

8.1 Principal-agent example

Consider again our principal-agent example, without external enforcement. In the underlying game, the worker chooses effort $a_1 \geq 0$, the outcome is $x = a_1$, and the payoff vector is $u(a_1) = (-a_1^2, a_1 + a_1^2)$.

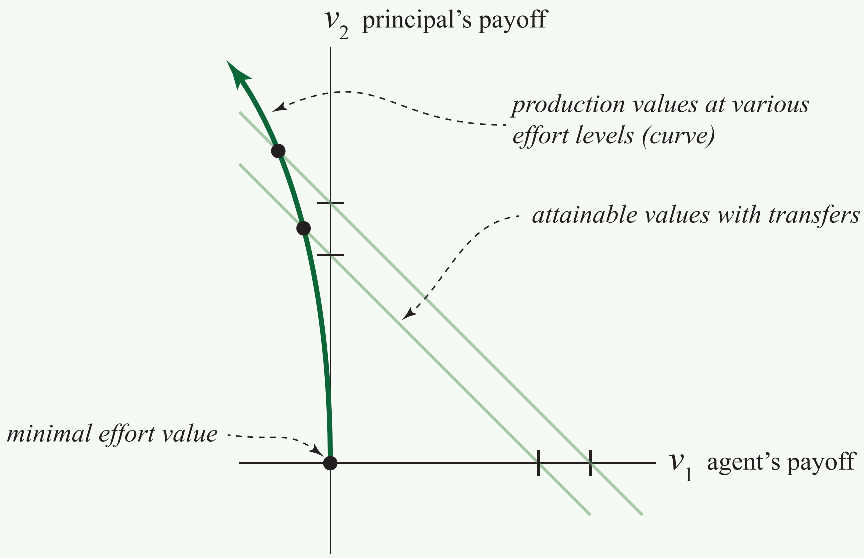
v_2 principal's payoff
production values at various effort levels (curve)
attainable values with transfers
minimal effort value
v_1 agent's payoff

PPE value set:

The highest effort level sustained in a PPE (SPE here) is the largest effort a_1 satisfying

$$(1-\delta)a_1^2 \leq \delta(z_1^2 - z_1^1),$$

where z^1 and z^2 are the upper-left and lower-right endpoints of the W^* triangle. This is $\overline{a}_1 = \sqrt{(z_1^2 - z_1^1)\delta/(1-\delta)}$.

Because repetition of the stage-game Nash equilibrium is a SPE, and the players can guarantee themselves at least 0, we have $z_1^1 = 0$. Also, it is clear that $z_1^2 = \overline{a}_1$. So we have

$$\overline{a}_1 = \delta/(1-\delta).$$

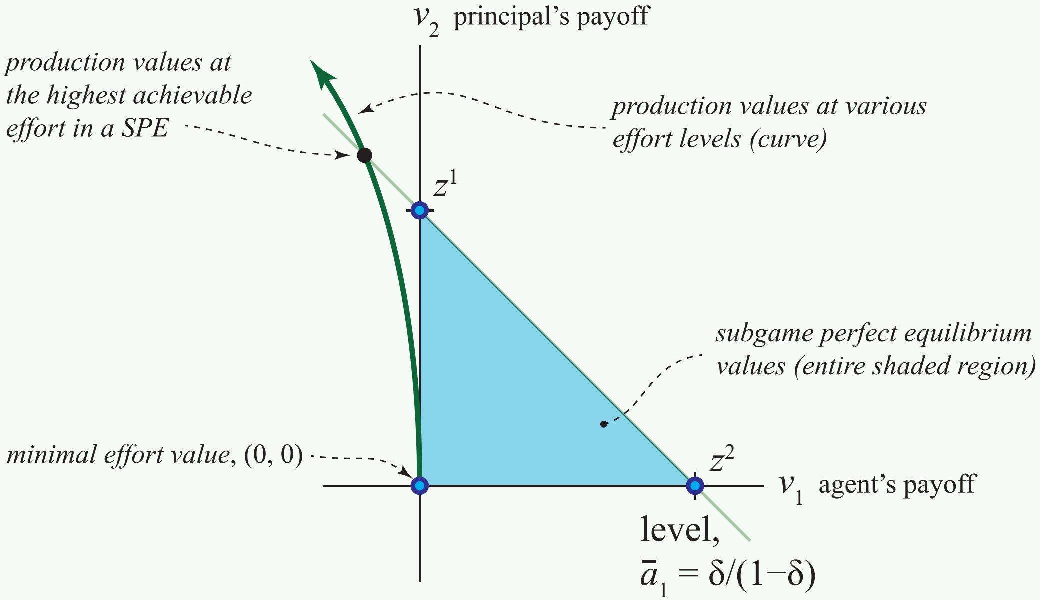

v_2 principal's payoff
production values at the highest achievable effort in a SPE
production values at various effort levels (curve)
z^1
subgame perfect equilibrium values (entire shaded region)
minimal effort value, (0, 0)
z^2
v_1 agent's payoff
level,
$\bar{a}_1 = \delta/(1-\delta)$

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 (for any history of play) is on the efficient frontier of W^*.

Goldlücke and Kranz (2013) provide general results for all settings with two players and perfect monitoring.

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 low high effort always.

8.2 Team-production example

The findings regarding Pareto perfection and pooled incentive conditions do not generally extend beyond the principal-agent setting, particularly to settings of imperfect monitoring in which more than one player takes an action in the stage game.

To see this and for another illustration of how to calculate PPE value sets, consider the team-production example we examined earlier, without external enforcement.

Underlying game: Players 1 and 2 simultaneously contribute effort. $A_1 = A_2 = \{0, 1\}$. High effort entails a personal cost of 4.

In the end, the joint project is either a success, which pays both players 10, or a failure, which pays them nothing. Success occurs with probability 0.8 if $a_1 = a_2 = 1$ (both players chose high effort), with probability 0.5 if $a_1 + a_2 = 1$ (exactly one of the players chose high effort), and with probability zero if $a_1 = a_2 = 0$.

The players do not observe each others' effort choices but they commonly observe whether the project is a success.

In the case of a failed project, a commonly observed noisy signal is generated about who is responsible. The signal identifies player 1 with probability $\frac{1}{2} - \frac{1}{2}\sigma(a_1 - a_2)$ and player 2 with probability $\frac{1}{2} - \frac{1}{2}\sigma(a_2 - a_1)$, where $\sigma \in [0, 1]$ is a parameter that measures the accuracy of the signal.

Note that if $\sigma = 0$ then the signal provides no information. If $\sigma = 1$ then it perfectly distinguishes between $a = (0, 1)$ and $a = (1, 0)$ but is still noisy in differentiating these action profiles from $a = (1, 1)$.

The outcome space is thus $X = \{1, 01, 02\}$, where $x = 1$ denotes a successful project, $x = 01$ denotes a failed project with player 1 identified as responsible, and $x = 02$ denotes failure with player 2 identified as responsible.

Recall that we write the underlying-game payoffs as a function of the action profile, averaging over the outcomes.

$$
\begin{aligned}
u(1,1) &= 0.8 \cdot (10,10) - (4,4) = (4,4) \\
u(0,1) &= 0.5 \cdot (10,10) - (0,4) = (5,1) \\
u(1,0) &= 0.5 \cdot (10,10) - (4,0) = (1,5) \\
u(0,0) &= 0 \cdot (10,10) - (0,0) \quad = (0,0).
\end{aligned}
$$

PPE value set:

The analysis of this example is similar to that for the previous example.. . .

The minimum payoff in the PPE value set is 1 for each player, so W^* is of the form

$$W^* = \text{conv}\{(1, L^* - 1), (L^* - 1, 1), (1, 1)\}.$$

The stage-game Nash equilibria $(0, 1)$ and $(1, 0)$ can be supported in the productive phase using a constant continuation-value function, which implies that $L^* \geq 6$.

The key issue is whether action profile $(1, 1)$ can be supported. For $(1, 1)$ to be a Nash equilibrium of the induced game, a deviator must be punished with a low continuation value in the event that the project fails.

Consider a candidate value set W with level L. There are two aspects of punishment: *individual* and *group*.

The best individual punishment gives the player who is identified as responsible the lowest individual continuation value, 1.

Suppose we do this by giving the other player the continuation value $L - 1 - k$, where k is an arbitrary number between 0 and $L - 2$.

- If we set $k = 0$, then this continuation value is on the frontier of W^*.
- However, if we set $k > 0$ then there is also an aspect of group punishment, where project failure reduces the joint value in the following period.

Group punishment helps provide effort incentives in the case of $\sigma < 1$ because a noisy signal will not perfectly identify a deviating player and, further, deviation by either player will increase the probability of failure (triggering the loss of joint value).

There is a trade-off here, because the reduction in joint value also occurs with positive probability when both players choose high effort, and thus it reduces equilibrium welfare.

To see the trade-off, write the incentive conditions for motivating high effort in the production phase.

- Contingent on a successful project, the players coordinate to achieve continuation value $(L/2, L/2)$ in the next period.
- Contingent on failure, they coordinate to achieve the continuation value that gives 1 to the responsible player and $L - 1 - k$ to the other player.

The players have the same incentive condition:

$$(1-\delta)\cdot 4+\delta\left[\tfrac{4}{5}\cdot\tfrac{L}{2}+\tfrac{1}{5}\left(\tfrac{1}{2}\cdot 1+\tfrac{1}{2}(L-1-k)\right)\right]$$
$$\geq (1-\delta)\cdot 5+\delta\left[\tfrac{1}{2}\cdot\tfrac{L}{2}+\tfrac{1}{2}\left(\tfrac{1+\sigma}{2}\cdot 1+\tfrac{1-\sigma}{2}(L-1-k)\right)\right].$$

The left side is the expected payoff in the induced game of choosing high effort when the other player also chooses high effort, whereas the right side is the expected payoff of deviating to low effort.

On both sides, the bracketed term is the expected continuation value from the start of the next period. The second term inside the brackets is the probability of failure times the resulting continuation value.

Case of $\sigma = 1$:

In this case, the signal of responsibility is perfectly accurate conditional on exactly one of the players providing low effort, group punishment does not improve incentives and it only lowers the equilibrium value.

- The maximal joint value is obtained with $k = 0$ and we have $L = 8$ if the incentive condition is satisfied.
- Plugging $L = 8$ into the incentive condition and rearranging terms yields $\delta \geq 2/5$.
- Therefore, if $\delta \geq 2/5$ then cooperation at the highest level can be sustained and $W^* = \text{conv}\{(1,7),(7,1),(1,1)\}$.
- Otherwise, only the Nash equilibria of the stage game can be sustained and we have $W^* = \text{conv}\{(1,5),(5,1),(1,1)\}$.

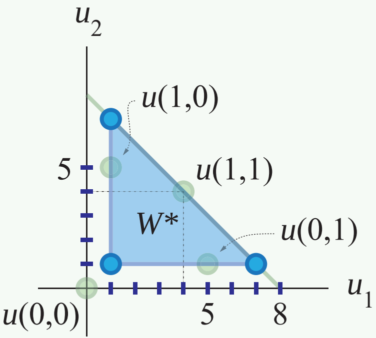

PPE value set, δ ≥ 2/5

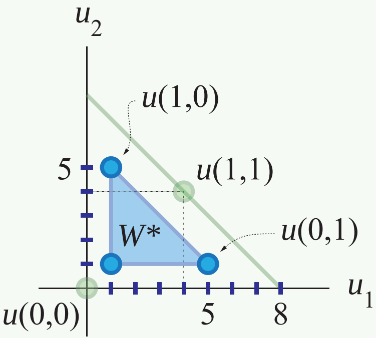

PPE value set, δ < 2/5

Case of $\sigma < 1$:

Here incentives may be still aligned without group punishment, achieving the maximum joint value, as long as players are patient enough.

- To find the cutoff discount factor, we plug $L = 8$ and $k = 0$ into the incentive condition and solve for δ, yielding $\delta \geq 2/(2+3\sigma)$.

- If $\delta \geq 2/(2+3\sigma)$ then the incentive condition can be relaxed by increasing the group punishment term k from zero. Cooperation featuring high effort by both players can be sustained but the level of the PPE set falls below 8.

- If the discount factor is below some number $\eta < 2/(2+3\sigma)$, then it becomes too costly to achieve high effort from both players and the PPE set is $W^* = \text{conv}\{(1,5),(5,1),(1,1)\}$.

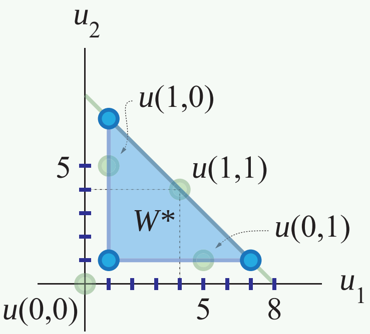

PPE value set, $\delta \geq 2/(2+3\sigma)$ PPE value set, $\eta > \delta \geq 2/(2+3\sigma)$ PPE value set, $\delta < \eta$

Note that when group punishments are needed, the equilibria that give the highest joint value entail continuation values on the equilibrium path (and off) that are in the interior of W^*.

Thus, these equilibria are not Pareto-perfect.

The logic behind group punishment extends what we found for external enforcement in the short-term setting to external enforcement in the relational setting.

In the short-term setting, we found implausible that the players would allow the external enforcer to follow through with a plan to extract money from them as a group in the event of project failure. The players would prefer to renegotiate their contract in this event, unraveling the incentive plan.

The same is true here.

9 Review of Relevant Bargaining Theory

Noncooperative game-theory models of negotiation…

Cooperative game-theory models of negotiation…

Nash bargaining solution predicts agreement on value w^* that solves

$$\max_{w \geq \underline{w}} \prod_{i \in \{1,2,\ldots,n\}} (w_i - \underline{w}_i)^{\pi_i}$$

over some bargaining set, where $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$ is the vector of nonnegative bargaining weights that sum to 1.

In settings with transferrable utility, $w^* = \underline{w} + \pi \left(L - \sum_{i \in \{1,2,\ldots,n\}} \underline{w}_i \right)$, where L is the maximum joint value available in bargaining set.

10 Settings with Trivial External Enforcement, Active Contracting

Next let us model negotiation by explicitly accounting for active contracting, where players exercise bargaining power in the process of reaching agreements.

In the present context, where there is no external enforcement, a contract is an agreement only about future behavior to be self-enforced. One can account for negotiation either noncooperatively or cooperatively.

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. They develop both a fully noncooperative model and a hybrid version in which negotiation is modeled cooperatively as the Nash bargaining solution with fixed bargaining weights.

We'll use the hybrid version.

We can think of the players, in the negotiation phase of any period, as bargaining over:

- An immediate transfer,
- The action profile they will play in the current period, and
- 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.

Bargaining theory:

In the hybrid model, the bargaining protocol is represented by an exogenous vector of bargaining weights $\pi = (\pi_1, \pi_2, \ldots, \pi_n)$ satisfying $\pi_i \geq 0$ and $\sum_{i=i}^{n} \pi_i = 1$, the same in every period.

Bargaining in the negotiation phase is resolved according to the generalized Nash bargaining solution: The players coordinate to achieve a continuation value that maximizes their joint value and, by making an immediate transfer, distribute the surplus according to their bargaining weights.

The surplus is relative to a *disagreement point*, whereby there is no immediate transfer and the players coordinate to achieve some achievable 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.

Characterization of the contractual equilibrium value (CEV) set:

Suppose a set W gives the continuation values from the negotiation phase of any period. 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 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

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

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 every 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. Under suitable technical conditions, contractual equilibrium exists and W^* is compact.

Below is an illustration of the CEV set in a two-player setting. We only need to keep track of the endpoints z^1 and z^2.

The *span* of W^*, denoted by $\mathrm{Span}(W^*)$ d^*, is defined as the horizontal (equivalently vertical) length of the CEV set; that is, $\mathrm{Span}(W^*) = z_1^2 - z_1^1 = z_2^1 - z_2^2$. The span figures prominently in the analysis of examples.

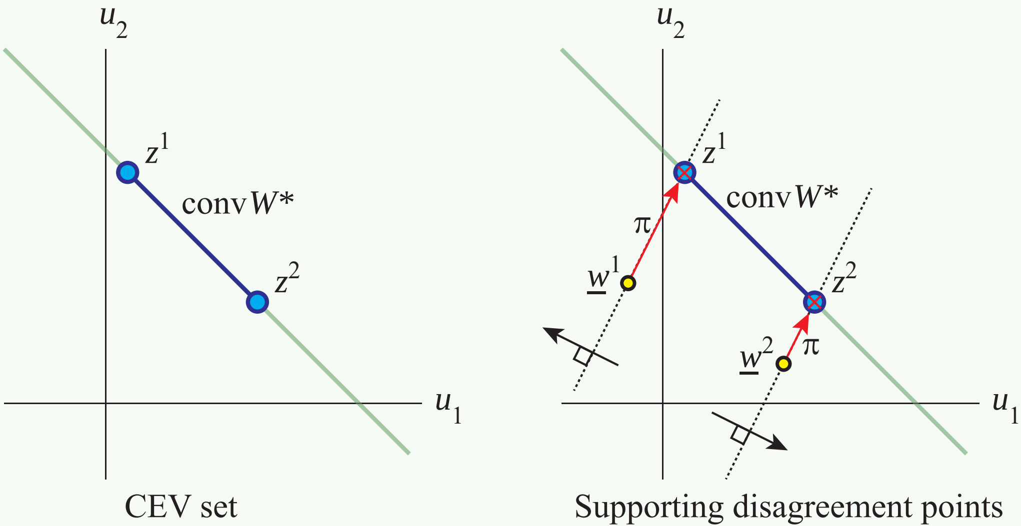
u_2
z^1
conv W*
z^2
u_1
CEV set
u_2
z^1
conv W*
π
$\underline{w}^1$
z^2
$\underline{w}^2$
π
u_1
Supporting disagreement points

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:

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

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 $(-\pi_2, \pi_1)$.

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 the figure above.

10.1 Principal-agent example

Consider again our principal-agent example. In the underlying game, the worker chooses effort $a_1 \geq 0$, the outcome is $x = a_1$, and the payoff vector is $u(a_1) = (-a_1^2, a_1 + a_1^2)$.

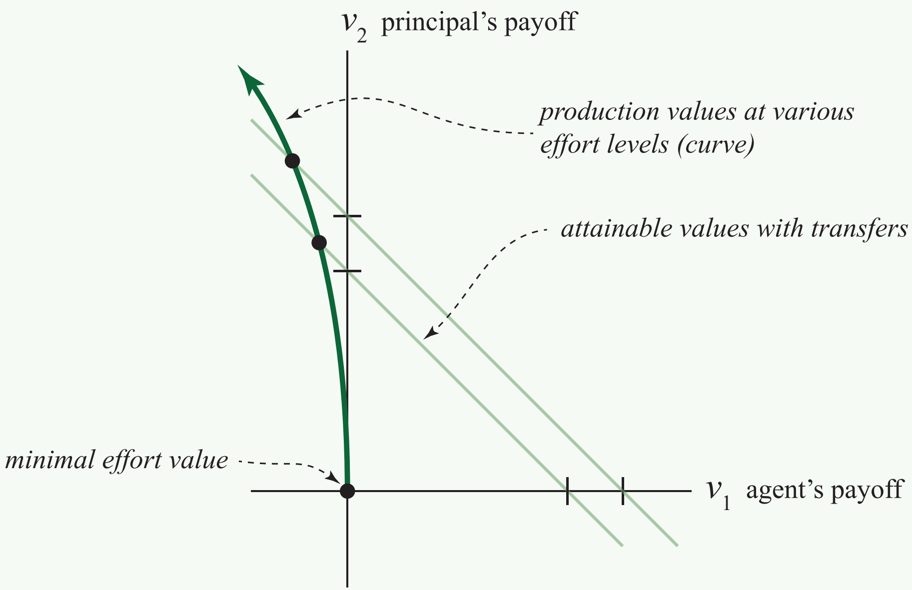
v_2 principal's payoff
production values at various effort levels (curve)
attainable values with transfers
minimal effort value
v_1 agent's payoff

As before, the highest effort level sustained in a CE is the largest effort a_1 satisfying

$$(1-\delta)a_1^2 \leq \delta\,\mathrm{Span}(W^*) = \delta(z_1^2 - z_1^1),$$

where z^1 and z^2 are the upper-left and lower-right endpoints of line segment W^*. This is $a_1^* = \sqrt{(z_1^2 - z_1^1)\delta/(1-\delta)}$. We have $L^* = a_1^*$

Determining z^1:

- The agent can guarantee herself at least 0.
- a_1^* can be enforced using the most sever punishment for any deviation, which is continuation value z^1.
- The agent's reward for cooperation can be set to make her indifferent, so that her continuation value from the current-period production phase is $0 + \delta z^1$. (It will turn out that the reward is to get z^2 from the next period.)
- This is the disagreement point $\underline{w}^1$ and it has joint value L^*.
- With no renegotiation surplus, $z^1 = \underline{w}^1$ and for the agent, $z_1^1 = \underline{w}_1^1 = 0 + \delta z^1$. Therefore $z^1 = (0, L^*) = (0, a_1^*)$.

Determining z^2:

- From the production phase, the principal's least favorite attainable continuation value is also the agent's most favorite. It involves zero effort in the current period and continuation value z^2 from the next period.
- Thus, $\underline{w}^2 = (0, 0) + \delta z^2$.
- The negotiation surplus is $(1 - \delta)(z_1^2 + z_2^2) = (1 - \delta)L^*$.
- So $z^2 = \underline{w}^2 + \pi(1 - \delta)L^* = \delta z^2 + \pi(1 - \delta)L^*$, implying $z^2 = \pi L^*$.

Determining $L^* = a_1^*$:

- We already evaluated the incentive condition for a_1^* to obtain

$$a_1^* = \sqrt{(z_1^2 - z_1^1)\delta/(1-\delta)}.$$

- Plugging in the values $z_1^1 = 0$ and $z_1^2 = \pi_1 a_1^*$ yields $a_1^* = \sqrt{\pi_1 a_1^* \delta/(1-\delta)}$.
- Solving, we get $a_1^* = \pi_1 \delta/(1-\delta)$.

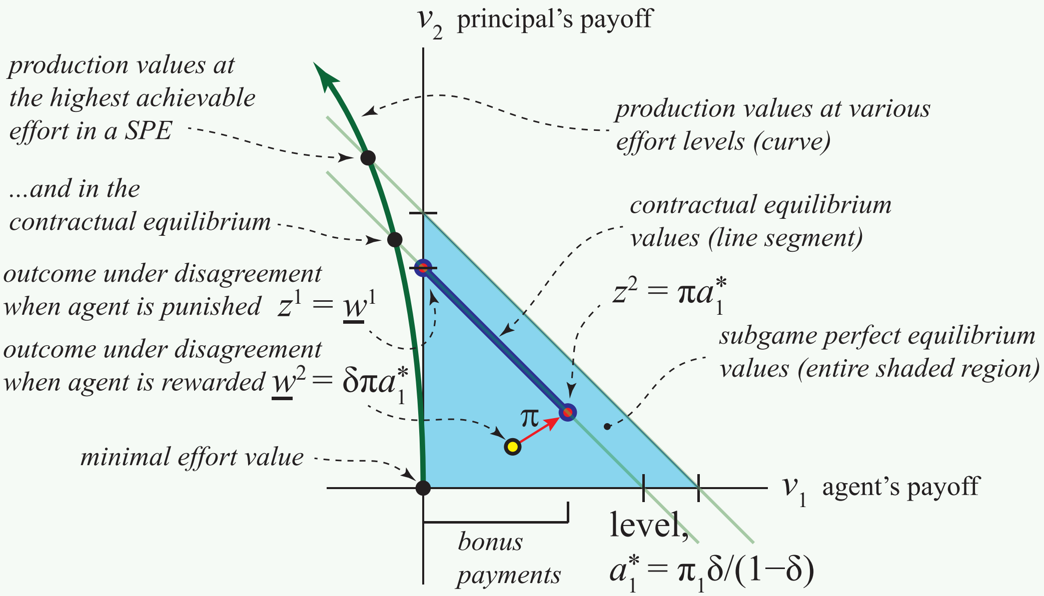
v_2 principal's payoff
production values at the highest achievable effort in a SPE
production values at various effort levels (curve)
...and in the contractual equilibrium
contractual equilibrium values (line segment)
outcome under disagreement when agent is punished $z^1 = \underline{w}^1$
$z^2 = \pi a_1^*$
subgame perfect equilibrium values (entire shaded region)
outcome under disagreement when agent is rewarded $\underline{w}^2 = \delta\pi a_1^*$
π
minimal effort value
v_1 agent's payoff
bonus payments
level, $a_1^* = \pi_1\delta/(1-\delta)$

Implications:

- CE predicts lower sustainable effort than does the best PPE in the inactive-contracting model: $a_1^* = \pi_1\delta/(1-\delta) < \delta/(1-\delta) = \overline{a}_1$ for $\pi_1 < 1$.

- Likewise, CE predicts a lower joint value L^* than one might predict in the inactive-contracting model.

- CE predicts a unique joint value, in contrast to a large range of PPE in the inactive-contracting model.

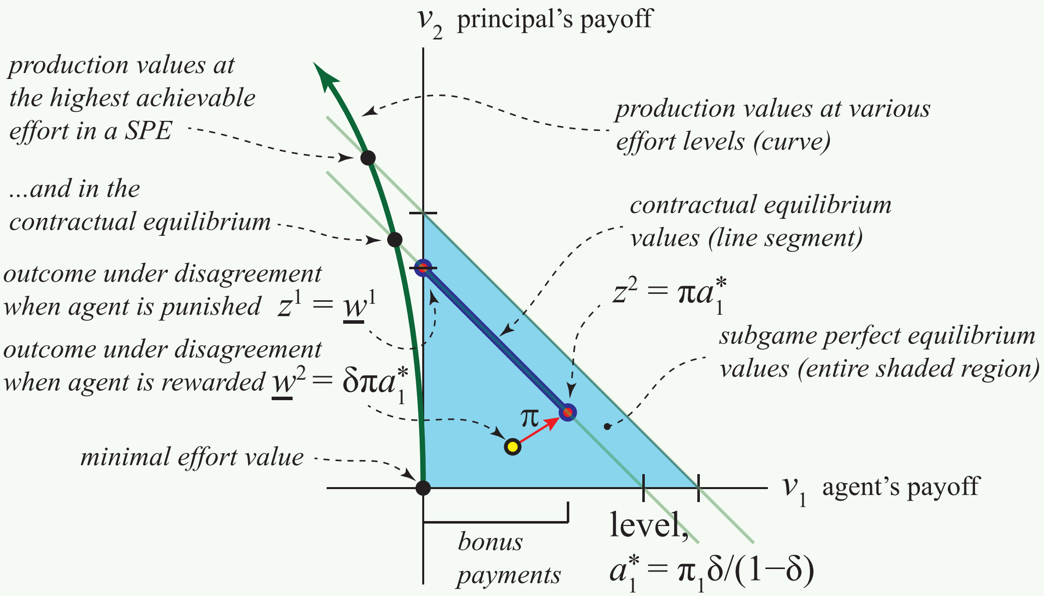

v_2 principal's payoff
production values at
the highest achievable
effort in a SPE
production values at various
effort levels (curve)
...and in the
contractual equilibrium
contractual equilibrium
values (line segment)
outcome under disagreement
when agent is punished $z^1 = \underline{w}^1$
$z^2 = \pi a_1^*$
subgame perfect equilibrium
values (entire shaded region)
outcome under disagreement
when agent is rewarded $\underline{w}^2 = \delta\pi a_1^*$
π
minimal effort value
v_1 agent's payoff
bonus
payments
level,
$a_1^* = \pi_1\delta/(1-\delta)$

10.2 Team-production example

For another illustration of contractual equilibrium, consider the team-production example without external enforcement.

Features of the CEV set:

- Renegotiation rules out group punishment, so sustaining high effort from both players relies on $\sigma > 0$.

- Disagreement point $\underline{w}^i$ entails play of player i's least preferred Nash equilibrium in the underlying game, followed by continuation value z^i, implying $\underline{w}^1 = (1-\delta)(1,5) + \delta z^1$ and $\underline{w}^2 = (1-\delta)(1,5) + \delta z^2$.

If $a = (1, 1)$ turns out not to be sustainable in equilibrium, then there is no bargaining surplus in these cases, implying $z^1 = (1, 5)$ and $z^2 = (5, 1)$.

If $a = (1,1)$ can be sustained in equilibrium, then $L^* = 8$ and the bargaining surplus ends up being $2(1-\delta)$, which implies $z^1 = (1,5) + 2\pi$ and $z^2 = (5,1) + 2\pi$.

- Player i's incentive condition for high effort, after some simplifying, is

$$(1-\delta) \cdot 4 + \delta \left[\tfrac{4}{5}\hat{w}_i + \tfrac{1}{5}(3 + 2\pi_i)\right] \geq (1-\delta) \cdot 5 + \delta \left[\tfrac{1}{2}\hat{w}_i + \tfrac{3}{2} + \pi_i - \sigma\right],$$

  where $\hat{w} \in W^*$ is the continuation value the players coordinate to achieve in the event of success.

- Recognizing that $\hat{w}_1 + \hat{w}_2 = 8$ and $\pi_1 + \pi_2 = 1$, when we can add the constraints for the two players, we get $10\sigma \geq (1-\delta)/\delta$.

- In fact, this constraint is necessary and sufficient for the two individual constraints because $\hat{w}$ can be set to balance slackness.

- High effort from both players is achieved if and only if $\delta \geq 1/(1+\sigma)$, a tighter condition than in the model with inactive contracting.

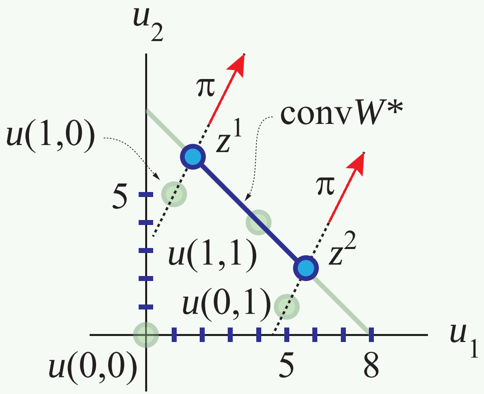

CEV set, δ ≥ 1/(1+σ)

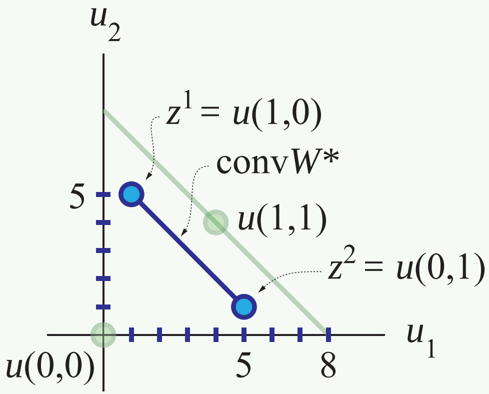

CEV set, δ < 1/(1+σ)

11 Settings with Nontrivial External Enforcement, Active Contracting

Let's now examine nonstationary settings in which b may vary over time due to nontrivial external enforcement and long-term external contracts.

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.

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.

11.1 Modeling ingredients

Drop the assumption made earlier that the transfer function b is exogenously fixed at $\mathbf{0}$.

Let $\mathcal{B} \equiv \{b\colon X \to \mathbb{R}^n_0\}$ 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 $c\colon H^X \to \mathcal{B}$, where for any $(t-1)$-period history $h \in H^X$, the transfer function specified for period t is $b^t = c(h)$.

Continuation contract:

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 $\tau - 1$.

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 history in which x is followed by the sequence h.

Define $c|x\colon H^X \to \mathcal{B}$ by $(c|x)(h) \equiv c(xh)$ for every $h \in H^X$.

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.

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.[3]

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

Active contracting in settings with two players:

- 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.

- The players enter any period t with continuation contract $\hat{c}^t$ in place from the previous period, they negotiate to change it to c^t and make transfer m^t.

- 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 the 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 settings studied previously.

11.2 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.

Continuation values 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

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

Recalling that the 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 $D(c(h^0), F^c(\mathcal{W}))$.

Continuation values from the negotiation phase with inherited continuation contract $\hat{c}$:

Next we apply the bargaining solution.

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 c-supported continuation value that maximizes their joint value,

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

and they choose an immediate transfer to achieve continuation value

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

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

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 rather complicated model may seem very difficult 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), simplifying the analysis, and helping to answer the questions just now posed.

The optimal continuation contract c^*, which achieves 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.
- 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 $2, 3, \ldots$., regardless of the history.
- Both on and off the equilibrium path, in each period the players renegotiate back to this same continuation contract.
- Thus, in period 2 the players 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.

Intuition behind this result:

- 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.
- 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.
- Future renegotiation will ensure that the high value is achieved in future periods as well, without reducing the span of continuation values.

Algorithm for calculating b^* and $\underline{b}$:

- Suppose we exogenously fix a single transfer function $\hat{b}$ for all periods, 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.[4]

[4] Other 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).

11.3 Project-choice example

The worker (player 1) chooses between three projects to apply one unit of effort, or nothing. The manager (player 2) observes the worker's effort choice and receives the revenue that it generates.

The manager has no action in the underlying game. The outcome includes a noisy binary signal of the worker's effort.

The set of feasible effort choices is $A_1 = \{0, 1, 2, 3\}$, where $a_1 = 0$ represents no effort, $a_1 = 1$ means applying effort to project 1, $a_1 = 2$ means applying effort to project 2, and $a_1 = 3$ means applying effort to project 3.

The signal is 1 with probability $\sigma(a_1)$ and 0 with probability $1 - \sigma(a_1)$. For each effort level, the following table gives the worker's effort cost, the manager's revenue, and the probability of the high signal:

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

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 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.[5] Let us assume that the signal is verifiable but player 1's effort choice is not verifiable, meaning that 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.

[5]Note that the contingent distribution function λ is given by $\lambda(0)(00) = 1$, $\lambda(1)(10) = 1/2$, $\lambda(1)(11) = 1/2$, $\lambda(2)(20) = 3/4$, $\lambda(2)(21) = 1/4$, $\lambda(3)(30) = 0$, and $\lambda(3)(31) = 1$.

The payoff vectors for the underlying game, along with the frontier of feasible values utilizing transfers, are pictured here:

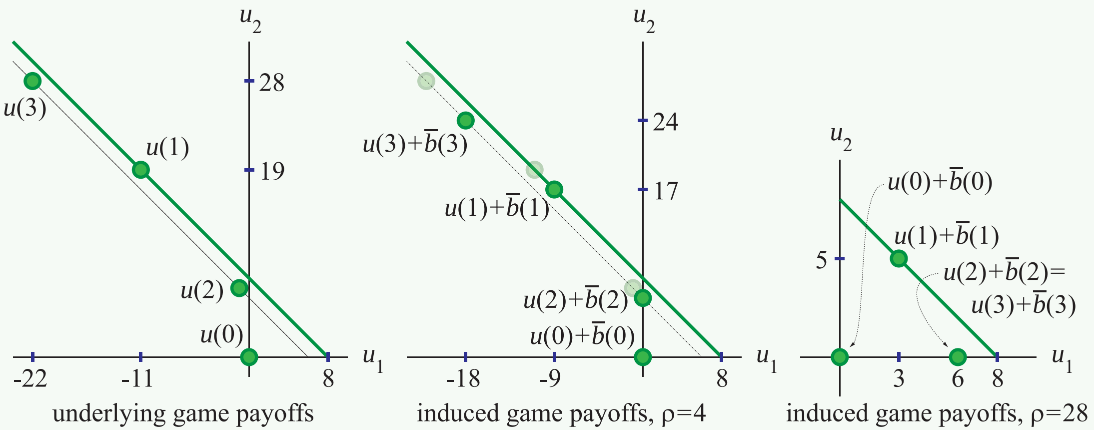

The middle and right graphs show the induced game for two different contracts, one specifying $\rho = 4$ and one specifying $\rho = 28$.

Results and implications:

- The players can achieve a joint value of 6 using only external enforcement, by agreeing to a bonus in each period that satisfies $\rho \geq 4$.

- The players cannot do better without intertemporal self-enforced rewards and punishments, because no contract can implement the efficient action $a_1 = 1$. This is easy to see by noting that $u_1(0) + \bar{b}(0) > u_1(1) + \bar{b}(1)$ for $\rho < 22$, $u_1(2) + \bar{b}(2) > u_1(1) + \bar{b}(1)$ for $\rho < 40$, and $u_1(3) + \bar{b}(3) > u_1(1) + \bar{b}(1)$ for $\rho > 22$.

- The difference between $\max\{u_1(a_1)+\bar{b}(a_1) \mid a_1 = 0, 2, 3\}$ and $u_1(1)+\bar{b}(1)$, though strictly positive, is minimized by choosing $\rho = 28$. 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.

- If the players' contract specifies $\rho = 4$, then the induced game has two Nash equilibria, $a_1 = 0$ and $a_1 = 2$.

- The optimal external contract is *semi-stationary*, specifying $\rho = 28$ in the current period and $\rho = 4$ in all future periods.

- In a CE, this is what the parties agree to, and in every period they renegotiate to select it again.

- Watson, Miller, and Olsen (2020) provide a general result along this line.

- They also prove a result establishing that external and self-enforcement are generally complementary (reversing conclusion of Baker, Gibbons, and Murphy 1994, 2002 and Schmidt and Schnitzer 1995).

11.4 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.[6]

An implication is that improving the external enforcement technology can have the effect of tightening incentive conditions and reducing welfare.

[6] A 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 inactive way (such as via a Nash-demand protocol as in Rayo 2007) or simply do not allow for renegotiation (Barron et. al. 2019 is an example).

In contrast, Watson, Miller, and Olsen (2020) show that, in the (perhaps more reaslistic?) setting in which players can renegotiate both components of their contract, the external-enforcement technology always complements self-enforcement.