

Relational Contracting, Repeated Negotiations, and Hold-Up*

Sebastian Kranz[†]

February, 2013

Abstract

We propose a unified framework to study relational contracting and hold-up problems in infinite horizon stochastic games. We first illustrate that with respect to long run decisions, the common formulation of relational contracts as Pareto-optimal public perfect equilibria is in stark contrast to fundamental assumptions of hold-up models. We develop a model in which relational contracts are repeatedly newly negotiated during relationships. Negotiations take place with positive probability and cause by-gones to be by-gones. Traditional relational contracting and hold-up formulations are nested as opposite corner cases. Allowing for intermediate cases yields very intuitive results and sheds light on many plausible trade-offs that do not arise in these corner cases. We establish a general existence result and a tractable characterization for stochastic games in which money can be transferred.

1 Introduction

In many economic relationships, parties can conduct investments, exert efforts or perform other actions that over shorter or longer time horizons determine their

*An earlier version of this paper had the title “Relational Contracting in Dynamic Stochastic Games: Repeated Negotiations and Hold-Up”

[†]University of Ulm, Department of Mathematics and Economics. Email: sebastian.kranz@uni-ulm.de. I would like to thank the German Research Foundation (DFG) for an individual research grant for visiting Yale University. This project started while I was visiting the Cowles Foundation in Yale and continued while I was working at the Institute for Energy Economics at the University of Cologne. I immensely benefited from the great hospitality and stimulating environment at both places. I would like to thank Mehmet Ekmekci, Eduardo Faingold, Susanne Goldlücke, Paul Heidhues, Johannes Hörner, David Miller, Larry Samuelson, Klaus Schmidt, Patrick Schmitz, Philipp Strack, Juuso Välimäki, Joel Watson and seminar participants at Arizona State University, UC San Diego, Yale, Bonn, Berlin and Munich for very helpful discussions.

joint surplus and possibly affect the way how that surplus is distributed. Limitations to formal contracting in economic relationships have inspired two large branches of economic literature that study relational contracts and hold-up problems, respectively.

Grout's (1984) classical article illustrates the essence of hold-up problems. He shows how firms under-invest in capital because labor unions appropriate a share of the generated surplus in subsequent wage negotiations.¹ The essence of relational contracting is to use repeated interactions and credible punishments to implement mutually desirable behavior.²

Despite the common motivation and economists' immense interest in both fields, a comprehensive framework for a unified analysis of relational contracting and hold-up problems is still missing. Relational contracts are typically formulated as Pareto-optimal public perfect equilibria (PPE) of infinitely repeated games.³ A limitation of repeated games is that players face the same stage game in every period, which restricts the ability to model relationships with long-term investments and corresponding hold-up problems. We study stochastic games, in which action spaces in each period depend on a state, which can change over time and be influenced by players' actions. They provide a natural framework for unified analysis of relational contracts, investments and associated hold-up problems.

While characterizations of Pareto-optimal PPE in stochastic games can quickly become intractable, scope for simplification arises from the fact that in most applications of relational contracting and hold-up problems, players have the opportunity to conduct monetary transfers with each other. The accompanying paper, Kranz (2012), shows that in stochastic games in which players can conduct voluntary monetary transfers, every PPE payoff can be implemented with a simple class of equilibria that have a stationary structure on the equilibrium path and use stick-and-carrot punishments. It also develops results that help to find Pareto-optimal PPE for any given discount factor.⁴

We will illustrate, however, that when relational contracting is augmented

¹Investment inefficiencies and the interaction with negotiation outcomes lie at the heart of hold-up problems. See Klein et. al. (1978), Williamson (1985) or Hart & Moore (1988) for other seminal contributions. See e.g. Schmitz (2002) for a survey.

²Examples for different applications of relational contracting include employment relationships by Levin (2002, 2003), MacLeod and Malcomson (1989) or Schmidt and Schnitzer (1995), the structure of firms by Baker, Gibbons & Murphy (2002) or Halonen (2002) and team production by Doornik (2006) and Rayo (2007), international trade agreements by Klimenko, Ramey and Watson (2008) or cartels by Harrington and Skrzypacz (2007, 2011). See Malcomson's (2010) for a survey.

³We use the term Pareto-optimal PPE to refer to a PPE that implements a payoff on the Pareto frontier of all PPE payoffs for a given discount factor. Even though for tractability reasons relational contracts are sometimes restricted to a simpler class of strategies, like grim-trigger strategies, the idea that players coordinate on Pareto-optimal equilibria essentially remains.

⁴These results extend the characterizations for repeated games with transfers by Levin (2003) and Goldlücke and Kranz (2012).

for long run decisions then *independent of the discount factor* many plausible hold-up problems are fully circumvented by Pareto-optimal PPE. The motivating example in Section 2 provides a simple illustration of this point with a classical hold-up example, in which players can conduct cooperative investments and no simple contractual solutions for the hold-up problem exist (see Che & Hausch, 1999). By flexibly coordinating trade decisions on the conducted investments, relational contracts in form of Pareto-optimal PPE always fully overcome the hold-up problem. An important insight is that incomplete formal contracting is not sufficient for the existence hold-up problems. Crucial is that also relational contracting is incomplete such that to a certain extend bygones are treated as bygones. Being able to account for this driving force of hold-up problems in relational contracting is the key motivation for introducing our concept of repeated negotiations of relational contracts.

An extreme form of incomplete relational contracting is given if each period earlier relational contracts are completely neglected and continuation play is always determined by new negotiations that ignore all payoff irrelevant aspects of history. This idea follows the spirit of the prevailing solution concept for stochastic games: Markov Perfect equilibria (MPE), in which only payoff relevant states can determine continuation play. If bygones are always bygones, hold-up problems fully reemerge. Yet, that assumption is orthogonally opposite to the essential feature of relational contracting: to coordinate continuation play in a flexible fashion on the history.

Our model allows a continuum of intermediate cases. We assume that an existing relational contract can depreciate at the beginning of a period with an exogenous *negotiation probability* and is then replaced by a new relational contract. Negotiations of new relational contracts follow a simple random dictator bargaining procedure in which bygones are bygones in the sense that the new relational contract does not condition on any payoff irrelevant aspects of the history. In a *repeated negotiation equilibrium* (RNE) all selected relational contracts must be incentive compatible, taking into account future negotiations, and maximize the expected payoff of the player who can select the new relational contract.

A larger negotiation probability reflects the assumption that history-independent bargaining power plays a stronger role in the relationship. In the corner case of a zero negotiation probability, the original relational contract always stays in place and corresponds to a Pareto-optimal PPE. If the negotiation probability is one and the game has a unique Markov perfect equilibrium, the RNE corresponds to that MPE.⁵ That relational contracts depreciate randomly, i.e. new negotiations are triggered by sun-spot events, has certain intuitive appeal and allows a simple way to formalize a continuous measure for the importance of history-independent

⁵One difficulty in interpreting MPE is that for sufficiently rich state spaces there often are multiple MPE, some of which do not well capture an intuitive notion of treating bygones as bygones. Section 2 exemplifies for such a case how for a negotiation probability of 1 the RNE corresponds to a MPE that particularly well captures the notion of bygones.

bargaining power.

From a theoretic perspective, repeated negotiation equilibria are complex objects. They form a fixed point of mutually optimal relational contracts, chosen by different players in different states. Relational contracts themselves form a PPE of a modified infinite horizon stochastic game that accounts for future negotiations. Given these complexities, there may be little hopes for a general existence result or a tractable characterization.

Indeed, we illustrate that equilibria do not generally exist if one assumes that players consider future negotiation outcomes as fixed when contemplating a deviation from a contract choice today. An important feature of RNE is that in no state a player prefers an alternative relational contract that would be profitable and incentive compatible under the belief that the player would choose that alternative contract again when future negotiations take place in the same state. We will illustrate that this belief follows naturally from the principle that by-gones shall be treated as by-gones in negotiations.

The main theoretical contribution of this paper is a general existence theorem for RNE. The theorem also shows that there always exist RNE with a tractable canonical form: all relational contracts constitute simple equilibria and negotiations affect the path of play only by changing the voluntary transfers that take place directly after the negotiations.

In infinite horizon games, repeated negotiations shed light on new, intuitive manifestations of hold-up problems that are closely interconnected with key features of relational contracting. Our main illustration is given by a classical relational contracting application, a principal-agent relationship, that is augmented for one long run decision: the principal can make herself permanently more vulnerable towards the agent by destroying her inside option.

From a traditional relational contracting perspective, increasing the own vulnerability is unambiguously beneficial as long as positive efforts by the agent are optimal on the equilibrium path. The principal will then destroy her inside option in all Pareto-optimal PPE. That is because a higher vulnerability allows for harsher punishments and thereby to implement higher effort on the equilibrium path. There is no drawback for the principal since Pareto-optimal PPE allow to perfectly coordinate away from any undesired abuse of the created vulnerabilities. In contrast, from a pure hold-up perspective, it is inadvisable to make oneself unilaterally more vulnerable, since it deteriorates the own bargaining position in future negotiations.

With a positive negotiation probability, the principal solves a natural trade-off between these two forces. Intermediate negotiation probabilities allow a simple analysis of the comparative statics of this trade-off, while the two corner cases of traditional relational contracting (Pareto-optimal optimal PPE) and pure hold-up (here, the unique MPE) essentially provide no insights in that respect.

Several more examples explore the interactions between relational contracting, hold-up and repeated negotiations. We show how a positive negotiation

probability extends the outside option principle to relational contracting, renders blackmailing incredible, or can induce costly arms races even when raising arms against other players involves costs but no direct gains.

In the special case of an infinitely repeated game, actions have no payoff-relevant long run effects and negotiation outcomes are therefore not affected by past decisions. As result, in repeated games our new framework remains mathematically equivalent to the traditional formulation of relational contracts as Pareto-optimal PPE. Negotiation can be interpreted as a restart of the relationship and a positive negotiation probability simply adjusts the discount factor downwards.

The assumption that negotiation occurs exogenously and with the same probability in every state provides a natural baseline case to illustrate the interaction between hold-up problems and relational contracts. We also show that the existence and characterization results extend to the case that negotiation probabilities differ between states. Section 6 illustrates that this result is quite powerful. By transforming the state space and adapting negotiation probabilities, one can, e.g., easily extend our main results to a model in which negotiations only take place if some players actively attempt to force negotiations.

We are only aware of a few papers that have studied the interaction of investments, hold-up and relational contracting. Baker Gibbons and Murphy (2002), Halonen (2002) and Blonsky and Spagnolo (2007) study the optimal allocation of property rights and optimal relational contracting in a repeated game with investments that always fully depreciate after one period. Ramey and Watson (1997) and Halac (2012) consider long-term investments but assume that investments take place only in the first period and afterward players always negotiate new relational contracts for the ensuing repeated game.⁶ Our results complement this literature by providing a framework that allows for much more flexible specifications of relationships with long run and short run decisions and negotiations of relational contracts.

The idea that relational contracts can be renegotiated during the relationship has been already explored in the literature on renegotiation-proofness in repeated games, e.g. Farrell and Maskin (1989), Bernheim and Ray (1989) or Asheim (1991). A key assumption in renegotiation-proofness concepts is that any player can block any renegotiation that makes her worse off than if the original relational contract stayed in place. In contrast, a key feature of repeated negotiation equilibria is that negotiations can make those players worse off whom the current relational contract grants higher continuation payoffs than the payoffs consistent with history-independent bargaining power. New negotiations in our model typically entail a redistribution of surplus from one player to another.

⁶Che and Sákovics (2004) and Pitchford and Snyder (2004) study hold-up problems in stochastic games with sequential investment decisions but they don't focus on relational contracting. Instead they assume that once investment stops, the resulting surplus is split via an enforceable contract.

While a positive negotiation probability can severely hamper the scope for co-operation, renegotiation-proofness does often not restrict the ability to implement Pareto-efficient PPE if monetary transfers are possible.⁷

The idea of modeling relational contracting as repeated negotiations is strongly inspired by Miller and Watson’s (2011) work on contract equilibria in repeated games. They assume that new negotiations take place in every period and consider a negotiation procedure with an explicit disagreement point that can depend on payoff irrelevant aspects of the history. The main factor by which negotiations reduce the flexibility of relational contracting, is that in periods of disagreement, players will not conduct transfers to each other. In our framework, the negotiation probability provides instead a continuous measure for the inflexibility of relational contracting. While Miller and Watson focus on implications of negotiations and disagreement in repeated games, our focus lies on the interaction of relational contract and hold-up problems in stochastic games.

The structure of the remaining paper is as follows. Section 2 motivates our concept using a classical two period hold-up model. Section 3 introduces stochastic games with transfers and reviews the characterization of Pareto-optimal PPE with simple equilibria. Section 4 introduces the general formulation and characterization of repeated negotiation equilibria. Section 5 illustrates the concept for several relational contracting examples with long run decisions. Section 6 discusses the extensions to state-dependent negotiation probabilities and endogenous negotiation. Two appendices contain proofs and additional results.

2 Motivating Example

This section motivates our concept with a classical two-period hold-up application. In period 1, a buyer and a seller, indexed by $i = 1, 2$, can each perform investments a_i from a compact set A_i . Investment costs for player i are given by a non-negative function $c_i(a_i)$. Investments determine, possibly stochastically, the state x in period 2, which determines production cost of the seller $k(x)$ and the valuation of the buyer $b(x)$. The total surplus from trade in period 2 is given by $S(x) = b(x) - k(x)$.

In period 2, a Nash demand game specifies whether trade takes place and how the surplus is split. Each player i announces simultaneously the share $d_i \in [0, 1]$ that she demands of the trade surplus. If $d_1 + d_2 \leq 1$ the distribution is feasible and each player i receives her share $d_i S(x)$; otherwise no trade takes place and players get outside payoffs of 0. Payoffs in the second period are discounted with a discount factor $\delta \in [0, 1)$.

⁷That is because monetary transfers often allow Pareto-efficient asymmetric punishments, see e.g. Levin, 2003 and Baliga and Evans, 2000. The results depend, however, on the exact assumptions on the timing and form of renegotiation-proofness. Kranz & Ohlendorf (2009) provide a comprehensive overview.

First best investments a^* maximize the sum of expected payoffs given that trade takes place whenever it is ex-post efficient:

$$a^* \in \arg \max_a E_x[\max\{\delta S(x), 0\}|a] - c_1(a_1) - c_2(a_2).$$

In the hold-up literature it is commonly assumed that surplus from trade is split according to the Nash bargaining solution, which in our example corresponds to an equal split of $S(x)$. It is then not in general possible to implement both first best investments and ex-post efficient trading decisions, i.e. a hold-up problem can arise.⁸ Note that the model allows for cooperative investments, i.e. the seller's investments can influence the buyer's valuation and vice versa. Che and Hausch (1999) show that in this case, the hold-up problem cannot generally be resolved with simple contractual solutions.⁹

The following result states the straightforward observation that if we remove the assumption that the surplus is split according to the Nash bargaining solution, a Pareto-optimal subgame perfect equilibrium can always fully mitigate the hold up problem.

Proposition 1. *The buyer-seller game has a Pareto-optimal subgame perfect equilibrium in which trading takes place and first best investments are conducted.*

The proof is simple. Assume that first best investments a^* are strictly positive for at least one player (otherwise the result is trivial). The straight line segment in Figure 1 (left) illustrates the Pareto frontier of subgame perfect continuation equilibria in period 2 given a state x with strictly positive surplus from trade. Consider strategies in which a player who has unilaterally deviated from a^* gets a continuation payoff of 0 in all states. If no player has unilaterally deviated, we pick continuation equilibria that split the surpluses $S(x)$ such that each player gets at least her cost c_i reimbursed. Since the expected discounted joint surplus under first best investments are larger than total investment costs, such a split of trade surplus always exists. Note that even after a deviation in period 1, all continuation payoffs are Pareto optimal, i.e. the equilibrium is consistent with a traditional renegotiation-proofness requirement.

The result simply makes use of the fact that the Nash demand game has a wide span of Pareto-efficient continuation payoffs in period 2 from which Pareto-optimal relational contracts can flexibly pick depending on the actually conducted

⁸For example, assume that only the seller has an investment opportunity, which costs 70 and generates deterministically a joint surplus of 100. Since the Nash bargaining solution gives the seller half of the surplus, i.e. 50, she would not recoup her investment costs, even though investments are socially efficient.

⁹In the case of non-cooperative investments, i.e. the seller's investments only influence production cost and the buyer's investments only influence her valuation, the hold-up problem can be effectively mitigated by writing simple option contracts in the first period (e.g. Nöldecke and Schmidt, 1995), or by having well structured compensation rules in civil law (e.g. Edlin & Reichelstein, 1996, or Ohlendorf, 2009).




Figure 1: Set of continuation payoffs in period 2. The thick line segment in the right figure illustrates the range of expected continuation payoffs that can be implemented for a negotiation probability of $\rho = 0.6$.

investments. While the Nash demand game has the non-compelling feature that players cannot continue bargaining after incompatible demands, there are many more sensible bargaining games that robustly yield the same Pareto-frontier of continuation payoffs. For example, Chatterjee and Samuelson (1990) show that in infinitely repeated simultaneous offer bargaining games every individual rational distribution of the trading surplus can be implemented, even when refining to (trembling-hand) perfect equilibria or to an even stronger notion of universal perfection. The famous exception is the alternative offer bargaining game by Rubinstein (1982), which uniquely implements the Nash bargaining outcome. However, that uniqueness result is not robust with respect to plausible modifications of the bargaining game. For example, Avery and Zemsky (1994) show that the availability of actions that can delay bargaining or destroy value restores the folk theorem. Since there are many natural ways to generate multiplicity and players could avoid the hold-up problem if they were allowed in an initial stage to choose a bargaining game with multiple equilibria, imposing a Rubinstein bargaining game seems similarly restrictive as imposing an equilibrium selection requirement in which continuation equilibria are given by the Nash bargaining solution.

Imposing the Nash bargaining solution corresponds to the idea that previous non-enforceable agreements on how to split trade surplus are considered as by-gones and ignored once investment costs are sunk. In contrast, Pareto-optimal relational contracts are built around the idea that past non-enforceable agreements always remain valid. The former idea constitutes the cornerstone of the hold-up literature, while the latter idea forms the corner stone of the relational contracting literature. Our model does not attempt to answer which of those ideas is more suitable, but rather provides a framework that unites both ideas by allowing a continuum of intermediate cases.¹⁰

In the example, a natural formulation of intermediate cases would be to require

¹⁰See Ellingsen, Tore and Robles (2002) and Tröger (2002) for evolutionary arguments on appropriate equilibrium selection. Ellingsen and Johannesson (2004) investigate hold-up problems experimentally. Their results support the view that intermediate cases are plausible.

that continuation equilibrium payoffs must lie on a line segment around the Nash bargaining solution whose span is a certain fraction of the span of the Pareto frontier of all SPE continuation payoffs. In Figure 2 (right), this is illustrated for a fraction of 0.4 by the thick line segment on the Pareto frontier.

Our formulation of randomly occurring repeated negotiations provides one implementation of such intermediate cases, which can be naturally extended to infinite horizon stochastic games. At the beginning of period 2, the existing relational contract will be replaced by a newly negotiated one with an exogenous negotiation probability $\rho \in [0, 1]$. If such negotiation takes place, bargaining follows a simple random dictator protocol: each player is chosen with equal probability to select the new relational contract (the general model allows players to have different bargaining weights). Bygones are then considered bygones and the chosen dictator selects a new relational contract that maximizes her continuation payoff. Hence, independent of conducted investments, player 1 will pick the contract that implements the right-most payoff from the set subgame perfect continuation payoffs and player 2 will select the top-most payoff. Thus, conditional that negotiation takes place, expected payoffs are equal to the Nash bargaining solution. With probability $1 - \rho$ the old relational contract remains valid, i.e. the terms of trade can then flexibly depend on the observed investments.¹¹

Consider the case that a player has deviated from required investments and is supposed to be punished by zero continuation payoffs in all states. Given the possibility of negotiation in period 2, that player is still able to guarantee herself an expected continuation payoff of

$$\frac{1}{2}\rho S(x)$$

in every state x with positive surplus. Hence, the span of expected continuation payoffs that can be implemented in state x is a fraction $1 - \rho$ of the span of the subgame perfect continuation payoffs. Figure 2 (right) thus shows the range of implementable expected payoffs for $\rho = 0.6$.

Example with simple functional form

For further illustration, assume player i can choose investments $a_i \in \{0, 1\}$ and investment costs simply are $c(a_i) = a_i$. The state x in period 2 is a deterministic function of investments and the resulting trade surplus shall be given by

$$S(x(a)) = \gamma(a_1 + a_2)$$

¹¹Since we consider risk-neutral players, only expected continuation payoffs will matter for players' incentives to deviate from a given relational contract. Hence, there is little disadvantage of specifying intermediate cases as probabilistic mixtures of extreme outcomes.

where $\gamma > 1$ is a measure of social desirability of investments. First best investment levels are

$$a^* = \begin{cases} (1, 1) & \text{if } \delta \geq \frac{1}{\gamma} \\ (0, 0) & \text{otherwise.} \end{cases}$$

To implement first best investments, it is optimal to split the trade surplus equally on the equilibrium path and to punish a player who deviates from required investments with a zero continuation payoff if the relational contract is not newly negotiated in period 2. Player i then has no incentive to deviate from investing $a_i = 1$ if and only if

$$-1 + \gamma\delta \geq \frac{1}{2}\rho\delta\gamma. \quad (1)$$

In line with Proposition 1, we find that absent repeated negotiation ($\rho = 0$) players always implement first best investments since the incentive constraint (1) then simplifies to the condition that positive investments are conducted in the first best solution: $\delta \geq \frac{1}{\gamma}$. Even though a lower discount factor tightens the incentive constraints for fixed investment levels, it does not affect the ability to implement first best investments. That is because a lower discount factor also makes high investments levels less desirable from a social perspective.¹²

In the limit case of no discounting $\delta \rightarrow 1$, the incentive constraint for implementing first best investments simplifies to

$$\rho \leq \frac{2(\gamma - 1)}{\gamma} \equiv \bar{\rho}.$$

The term $\bar{\rho}$ denotes a critical negotiation probability above which it is not possible to implement first best investments. Similar to the common practice in repeated games to use critical discount factors, one can use the critical negotiation probability to conduct comparative statics of the players' ability to implement efficient long-run decisions. In our example, the comparative statics are not surprising: the critical negotiation probability decreases in the parameter γ that determines the surplus of investments. In dynamic stochastic games with long run decisions, critical negotiation probabilities have the conceptual advantage over critical discount factors that the first best decisions are not affected by the negotiation probability.

¹²For infinite horizon games, the following intuition will generally be useful. A reduction of the discount factor has different effects on the ability to implement first-best short-run and long-run actions, respectively. While implementation of first-best short-run actions generally becomes harder, the effect on first-best long-run actions is ambiguous since a lower discount factor reduces the social desirability of current costs compared to future benefits. In contrast, an increase in the negotiation probability does not change the first best solutions and symmetrically reduces the ability to implement first best long- and short-run actions.

Bygones and Markov Perfect equilibria

We conclude the motivating example with an observation on the relationship between bygones, negotiation in every period, and Markov perfect equilibria. In a Markov perfect equilibrium, continuation play in period 2 is only allowed to depend on the state x . Yet, Markov perfection does not restrict the ability to implement first best investments in our (functional form) example since the state x is sufficiently informative about the investment decisions. The requirement of Markov perfection does not imply a strong notion of bygones. Yet, the repeated negotiation equilibria for the case $\rho = 1$ is equivalent to a specific MPE that corresponds to a strong notion of bygones.

3 Stochastic Games with Transfers and Simple Equilibria

This section defines infinite horizon stochastic games with transfers and summarizes the key results in Kranz (2012) that show how every PPE equilibrium payoff can be implemented with a simple class of equilibria.

3.1 Stochastic Games with Transfers

We consider n -player stochastic games of the following form. There are infinitely many periods and future payoffs are discounted with a common discount factor $\delta \in [0, 1)$. There is a finite set of states X and $x^0 \in X$ denotes the initial state. A period is comprised of two stages: a transfer stage and an action stage. There is no discounting between stages. At the beginning of each period players commonly observe a public signal from a continuous distribution, which determines whether negotiations take place and which player can choose the new relational contract.

In the transfer stage, every player simultaneously chooses a non-negative vector of transfers to all other players.¹³ Players also have the option to transfer money to a non-involved third party, which has the same effect as burning money. Transfers are perfectly observed by all players.

In the action stage, players simultaneously choose actions. In state $x \in X$, player i can choose a pure action a_i from a finite or compact action set $A_i(x)$. The set of pure action profiles in state x is denoted by $A(x) = A_1(x) \times \dots \times A_n(x)$.

After actions have been conducted, a signal y from a finite signal space Y and a new state $x' \in X$ are drawn by nature and commonly observed by all players. We denote by $\phi(y, x'|x, a)$ the probability that signal y and state x' are drawn; it depends only on the current state x and the chosen action profile a . Player i 's

¹³To have a compact strategy space, we assume that a player's transfers cannot exceed an upper bound of $\frac{\delta}{1-\delta} \sum_{i=1}^n [\max_{x \in X, a \in A(x)} \pi_i(a, x) - \min_{x \in X, a \in A(x)} \pi_i(a, x)]$. That bound is large enough to be never binding given the incentive constraints of voluntary transfers.

stage game payoff is denoted by $\hat{\pi}_i(a_i, y, x)$ and depends on the signal y , player i 's action a_i and the initial state x . We denote by $\pi_i(a, x)$ player i 's expected stage game payoff in state x if action profile a is played. If the action space in state x is compact then stage game payoffs and the probability distribution of signals and new states shall be continuous in the action profile a .

We assume that players are risk-neutral and that payoffs are additively separable in the stage game payoff and money. This means that the expected payoff of player i in a period with state x , in which she makes a net transfer of p_i and action profile a has been played, is given by $\pi_i(a, x) - p_i$.

When referring to (continuation) payoffs of the dynamic stochastic game, we mean expected average discounted continuation payoffs, i.e. the expected sum of continuation payoffs multiplied by $(1 - \delta)$.

We either restrict attention to pure strategies or, for finite action spaces, also consider strategies in which players can mix over actions. If equilibria with mixed actions are considered, $\mathcal{A}(x)$ shall denote the set of mixed action profiles at the action stage in state x otherwise $\mathcal{A}(x) = A(x)$ shall denote the set of pure action profiles. For a mixed action profile $\alpha \in \mathcal{A}(x)$, we denote by $\pi_i(\alpha, x)$ player i 's expected stage game payoff taking expectations over mixing probabilities and signal realizations.

A public history describes the sequence of all states, public signals and monetary transfers that have occurred before a given point in time. A public strategy σ_i of player i in the stochastic game maps every public history that ends before the action stage into a possibly mixed action $\alpha_i \in \mathcal{A}_i(x)$, and every public history that ends before a payment stage into a vector of monetary transfers. A public perfect equilibrium (PPE) is a profile of public strategies that constitutes mutual best replies after every history. If actions can be perfectly monitored, i.e. $y = a$, PPE are equivalent to subgame perfect equilibria.

3.2 Simple Equilibria

A simple strategy profile is characterized by $n + 2$ phases. Play starts in the up-front transfer phase, in which players are required to make up-front transfers described by a vector of net payments p^0 .¹⁴ Afterwards play can be either in the *equilibrium phase*, indexed by $k = i$, or in the *punishment phase* of some player i , indexed by $k = i$.

A simple strategy profile specifies for each phase $k \in \mathcal{K} = \{e, 1, \dots, n\}$ and state x an action profile $\alpha^k(x) \in \mathcal{A}(x)$. We refer to α^e as the equilibrium phase policy and to α^i as the punishment policy for player i and call the collection of all policies $(\alpha^k)_{k \in \mathcal{K}}$ a policy plan. From period 2 onwards, required net transfers are described by a vector $p^k(x', y, x)$ that depends on the current phase k , the current state x' , and the realized signal y and state x of the previous period.

¹⁴In a simple equilibrium transfers will always be structured such that no player at the same time makes transfers and receives transfers.

The transitions between phases are simple. If no player unilaterally deviates from a required transfer, play transits to the equilibrium phase: $k = e$. If player i unilaterally deviates from a required transfer, play transits to the punishment phase of player i , i.e. $k = i$. In all other situations the phase does not change.

This means that punishments in a simple equilibrium have a stick and carrot structure. They never last longer than one period and will be settled by a transfer of the punished player that can be interpreted as payment of a fine to the other players. Like in a Markov perfect equilibrium, actions on the equilibrium path only depend on the state. Transfers on the equilibrium path are used to balance incentive constraints among different players, while upfront transfers allow a flexible distribution of joint equilibrium payoffs.

Let $\bar{U}(x)$ denote the supremum of the joint PPE continuation payoffs at the beginning of a period in state x and $\bar{v}_i(x)$ the corresponding infimum of player i 's PPE continuation payoffs. An *optimal simple equilibrium* shall be a simple equilibrium that implements in every state x in the equilibrium phase a joint payoff of $\bar{U}(x)$ and gives each player i in her punishment state a continuation payoff of $\bar{v}_i(x)$.

Theorem 1. (*Kranz, 2012*) *A stochastic game with voluntary transfers has an optimal simple equilibrium and by adjusting incentive compatible upfront transfers it can implement every PPE payoff. The set of PPE continuation payoffs in state x is closed and given by the simplex*

$$\{u \in \mathbb{R}^n \mid \sum u_i \leq \bar{U}(x) \text{ and } u_i \geq \bar{v}_i(x) \forall i\} \quad (2)$$

where $\bar{U}(x)$ denotes the maximum of joint payoffs and $\bar{v}_i(x)$ the minimum of player i 's payoffs across all PPE starting in state x .

Characterizing the set of PPE continuation payoffs boils down to finding $n + 1$ numbers for each state x : the highest joint payoffs $\bar{U}(x)$ and the lowest payoffs $\bar{v}_i(x)$ for each player i . Furthermore, we can restrict attention to finding an optimal simple equilibrium for being able to implement any desired PPE payoff. Kranz (2012) contains results for finding optimal simple equilibria. Figure 2 illustrates the PPE continuation payoff set for a two player game.

Similar to the Nash demand game studied in Section 2, the Pareto-frontier of PPE continuation payoffs is always linear.

4 Repeated Negotiation Equilibria

This section formulates and characterizes repeated negotiation equilibria for stochastic games with transfers.




Figure 2: Set of public perfect continuation equilibrium payoffs at the beginning of period in state x in a two player stochastic game with transfers.

4.1 Key concepts of repeated negotiations

A relational contract shall be an incomplete strategy profile that describes play just until new negotiations take place. We assume that at the beginning of a period a sunspot signal is observed with a negotiation probability $\rho \in [0, 1]$, which indicates that new negotiations take place. In the first period of the game negotiations always take place.

Negotiations shall follow a simple random dictator procedure: one player is randomly chosen to select the new relational contract. The probability that player i is chosen is denoted by β_i and called i 's bargaining weight. The selected relational contract shall only depend on the current state x and on the identity of the player that selects it. Furthermore, relational contracts shall not condition on any event that occurred before they were negotiated. A helpful picture is that players forget their history when negotiations take place, i.e. payoff irrelevant aspects of the history are then completely treated as bygones.

We denote by $\sigma_{(i,x)}$ the relational contract selected by player i in state x . A profile of selected relational contracts for all states and players

$$\sigma = \times_{i \in \{1, \dots, n\}, x \in X} \sigma_{(i,x)}$$

is called a contract profile. Every contract profile constitutes a strategy profile of the stochastic game, in which the sun spot signal at the beginning of a period specifies whether negotiations take place and who selects the new relational contract. We denote by $\sigma_{-(i,x)}$ a contract profile that excludes the relational contract selected by player i in state x .

For a given contract profile, we denote by $r_j^i(x|\sigma)$ player j 's continuation payoff in the stochastic game directly after negotiations have taken place in state x and player i has selected the new relational contract $\sigma_{(i,x)}$. Generally, we refer

to a function \mathbf{r} that maps every pair (i, x) of player and state into a payoff vector (bounded by the range of feasible payoffs) as negotiation payoffs and denote by \mathcal{R} the set of negotiation payoffs; $\mathbf{r}(\cdot|\sigma)$ are the negotiation payoffs of the contract profile σ .

Truncated games Before defining repeated negotiation equilibria, we introduce a class of truncated games, which provide a convenient tool to analyze negotiation payoffs and to determine incentive compatibility of relational contracts taking account of future negotiations. A truncated game $\Gamma(\mathbf{r}, x_s)$ is parametrized by arbitrary negotiation payoffs \mathbf{r} and an initial state $x_s \in X$. As long as no new negotiation has taken place, payoffs and action spaces of the truncated game are the same as in the original game. If negotiations take place in state x and player i chooses the new contract, play transits to an absorbing state in which players automatically get fixed payoffs $\mathbf{r}^i(x)$ in every future period, i.e. the truncated game essentially ends. The truncated game has no negotiation in the first period, i.e. there is at least one period of play before an absorbing state is reached.

The definitions directly imply

Lemma 1. *A contract profile σ constitutes a PPE of the original game if and only if for every player i and every state x , the relational contract $\sigma_{(i,x)}$ constitutes a PPE of the truncated game $\Gamma(\mathbf{r}(\cdot|\sigma), x)$.*

For a given contract profile σ and some negotiation payoffs \mathbf{r} , let $\mathbf{g}^i(x|\sigma, \mathbf{r})$ denote the payoffs of the relational contract $\sigma_{(i,x)}$ in the truncated game $\Gamma(\mathbf{r}, x)$. We specify by $G^\sigma : \mathcal{R} \rightarrow \mathcal{R}$ an operator that maps negotiation payoffs into the payoffs of the corresponding truncated games, i. e.

$$G^\sigma(\mathbf{r}) \equiv \mathbf{g}(\cdot|\mathbf{r}, \sigma).$$

Let $d_\infty : \mathcal{R} \times \mathcal{R} \rightarrow \mathbb{R}_0^+$ be the metric induced by the supremum norm. The next result establishes a useful link between the negotiation payoffs of a contract profile and the payoffs of its individual relational contracts in the corresponding truncated games.

Lemma 2. *G^σ is monotone increasing and, in the metric space (\mathcal{R}, d_∞) , a contraction mapping that has a unique fixed point given by the negotiation payoffs induced by σ in the original game. Therefore*

$$\mathbf{r} = \mathbf{g}(\cdot|\mathbf{r}, \sigma) \Leftrightarrow \mathbf{r} = \mathbf{r}(\cdot|\sigma).$$

4.2 Repeated Negotiation Equilibria

Consider a relational contract $\tilde{\sigma}_{(i,x)}$ chosen by player i in state x and take as given a profile of other contracts $\sigma_{-(i,x)}$. We say $\tilde{\sigma}_{(i,x)}$ is incentive compatible if $\tilde{\sigma}_{(i,x)}$ constitutes a PPE of the truncated game $\Gamma(\mathbf{r}(\cdot|\tilde{\sigma}_{(i,x)}, \sigma_{-(i,x)}), x)$, i.e. no player

shall have an incentive to deviate from $\tilde{\sigma}_{(i,x)}$ if player i always selects it in state x . We say $\tilde{\sigma}_{(i,x)}$ is *strictly preferred* over another relational contract $\hat{\sigma}_{(i,x)}$ if

$$r_i^i(x|\tilde{\sigma}_{(i,x)}, \sigma_{-(i,x)}) > r_i^i(x|\hat{\sigma}_{(i,x)}, \sigma_{-(i,x)}),$$

i.e. if always chosen in state x , it grants player i a larger negotiation payoff.

Definition 1. A contract profile σ constitutes a *negotiation equilibrium (RNE)* if for every state x and every player i , the relational contract $\sigma_{(i,x)}$ is incentive compatible and there exists no relational contract $\tilde{\sigma}_{(i,x)}$ that is also incentive compatible given $\sigma_{-(i,x)}$ and is strictly preferred over $\sigma_{(i,x)}$.

An important element of the definition is that, loosely speaking, selecting an alternative relational contract is not treated as a one shot deviation: If today player i selects an alternative relational contract $\tilde{\sigma}_{(i,x)} \neq \sigma_{(i,x)}$, the incentive compatibility and profitability of the alternative contract is assessed under the belief that also in the future player i will select $\tilde{\sigma}_{(i,x)}$ in state x .

This assumption is a natural consequence of our notion that in negotiations by-gones are by-gones, which we symbolized by the picture that players forget the whole history of play when negotiations take place. If today in state x there are any reasons for why player i prefers to select the relational contract $\tilde{\sigma}_{(i,x)}$ and that contract is deemed incentive compatible, players should then rationally predict that the same reasons apply every time player i can select a relational contract in state x because the situation in the future will be exactly the same as today.

An equilibrium concept in which players would not anticipate that profitable deviations from contract choice would be repeated in future negotiations, would be plagued by non-existence problems. The black-mailing game in Section 5.1 will provide a simple illustration for this point.

4.3 Canonical Repeated Negotiation Equilibria and Existence

We say that σ is an incentive compatible *canonical contract profile* if all its relational contracts only differ by their upfront payments and constitute optimal simple equilibria of the truncated games with negotiation payoffs $r(\cdot|\sigma)$; if σ is also a RNE, we call it a canonical RNE. Based on Theorem 1 and the fixed point result in Lemma 2, we can establish

Proposition 2. *For every RNE σ there exists a incentive compatible canonical contract profile $\tilde{\sigma}$ that has the same negotiation payoffs.*

We cannot generally show, however, that for every RNE there also exists a canonical RNE that has the same negotiation payoffs. The problem is that if for

player i in state x , one substitutes the original relational contract with an optimal simple equilibrium that has the same payoffs, the substitution might enlarge the set of incentive compatible relational contracts for other players or in other states and potentially destroy optimality of some of the current contract choices (even though all current contracts will remain incentive compatible). Appendix A provides a sufficient condition on state transitions that rules out this possibility and ensures that the negotiation payoffs of every RNE can be implemented with a canonical RNE. That sufficient condition is satisfied by all examples in this paper.

We now state our general existence result for canonical renegotiation equilibria.

Theorem 2. *If the action space is finite and mixed actions are allowed then a canonical RNE exists.*

The proof of Theorem 2 is relatively involved and relies on a series of preliminary results. We refer the reader to Appendix A for a detailed development.

4.4 Regular Negotiation Payoffs

On first thought, it seems intuitive that in a RNE σ with negotiation payoffs \mathbf{r} , player i selects in state x a relational contract that grants her the highest PPE payoff of the truncated game $\Gamma(\mathbf{r}, x)$. In that case, we say that the relational contract $\sigma_{(i,x)}$ has *regular* negotiation payoffs, which satisfy

$$\begin{aligned} \mathbf{r}_i^i(x) &= \bar{U}(x|\mathbf{r}) - \sum_{j \neq i} \bar{v}_j(x|\mathbf{r}), \\ \mathbf{r}_j^i(x) &= \bar{v}_j(x|\mathbf{r}), \end{aligned}$$

where $\bar{U}(x|\mathbf{r})$ denotes the highest joint payoff and $\bar{v}_j(x|\mathbf{r})$ the lowest payoff of player j across all PPE payoffs of the truncated game $\Gamma(\mathbf{r}, x)$ (compare with Figure 2).

For given negotiation payoffs \mathbf{r} we denote by r (in non-bold fonts) the expected negotiation payoffs assuming that it is not yet known, which player can make the offer, i.e.

$$r(x) \equiv \sum_{i=1}^n \beta_i \mathbf{r}^i(x).$$

Expected regular negotiation payoffs satisfy

$$r_i(x) = \bar{v}_i(x|\mathbf{r}) + \beta_i (\bar{U}(x|\mathbf{r}) - \sum_{j=1}^n \bar{v}_j(x|\mathbf{r})).$$

They split the highest joint continuation payoff according to a generalized Nash bargaining solution in which the threat point is given by the profile of the lowest PPE payoffs for every player.

Even though in many examples, RNE have regular negotiation payoffs, this is not always the case, as the blackmailing game in the next section will illustrate.

5 Examples

This section illustrates the effects of repeated negotiation in relational contracting with simple examples. For the sake of clarity and simplicity all examples consider games in which actions can be perfectly monitored and restrict attention to pure strategy equilibria, i.e. every PPE is a subgame perfect equilibrium (SPE). Appendix B contains proofs of the results in this section.

5.1 The blackmailing game

We first consider a simple game to illustrate that RNE can have irregular negotiation payoffs. Player 1 (the blackmailer) has evidence about some illegal activity of player 2 (the target) and can decide in the initial state x_0 whether to reveal it $a = a_R$ or to keep it secret $a = a_S$. As long as the evidence has not been revealed, the state stays x_0 and once the evidence has been revealed, the game permanently moves to an absorbing state x_1 in which no more actions can be taken. Stage game payoffs are

$$\begin{aligned}\pi(a_S, x_0) &= (0, 1) \\ \pi(x_1) = \pi(a_R, x_0) &= (0, 0).\end{aligned}$$

Revealing the evidence involves no cost for the blackmailer but reduces the target's payoffs by 1 in the current and all future periods.

Consider a simple strategy profile in which the blackmailer reveals the evidence (only) if he punishes the target in state x_0 (for not having paid a specified bribe in the transfer stage). Regular expected negotiation payoffs would then be given by

$$\begin{aligned}\tilde{r}_1(x_0) &= \beta_1 \\ \tilde{r}_2(x_0) &= -\beta_1.\end{aligned}$$

Regular negotiation payoffs seem intuitive on first sight: the blackmailer extracts from the target an amount equal to the blackmailer's bargaining weight β_1 multiplied by the damage (measured in money) that is imposed on the target by revealing the evidence.

However, simple arguments show that in every RNE, the blackmailer must have in state x_0 a irregular expected negotiation payoff of zero. In state x_1 continuation payoffs are zero for both players. This implies that if and only if the blackmailer has zero expected negotiation payoffs in state x_0 , the truncated game $\Gamma(x_0, \mathbf{r})$ has a subgame perfect equilibrium in which the blackmailer reveals the evidence. That is because under a positive negotiation payoff the blackmailer would strictly prefer to stay in state x_0 . Having pinned down the blackmailer's negotiation payoffs, we can conclude that there is a RNE in which both players decide to neither conduct transfers nor to reveal the evidence.

Intuitively, one can interpret this RNE as the limit case of the following relational contracts. The target agrees to pay the blackmailer a very small amount $\varepsilon > 0$ for not revealing the evidence. Since renegotiation outcomes only depend on the state, both players know that when renegotiation takes place again, the blackmailer can again extort an amount of ε -magnitude from the target. Since any positive ε removes the blackmailer's incentives to reveal the evidence, the RNE must correspond to the limit case of $\varepsilon = 0$. While that result may seem surprising on first sight, it seems intuitive given that the blackmailer has no commitment device that prevents future extortion of the target.¹⁵

Since the blackmailer always gets a payoff of zero, there exist additional RNE in which the blackmailer selects a relational contract in which he reveals the evidence with positive probability or forces the target to burn money. Given the interpretation above, the Pareto-optimal RNE seems more plausible in this example, however.

Recall from Section 4.2 that the incentive compatibility of a relational contract in state x is assessed under the common belief that when player i selects an alternative relational contract today then player i will make the same decision whenever she selects again a relational contract in state x . A natural alternative formulation, would have been to hold future negotiation payoffs fixed when a different relational contract is chosen today. It is simple to see, however, that an equilibrium defined according to that alternative formulation would fail to exist in the blackmailing game. Whenever the blackmailer's negotiation payoff in state x_0 is zero, she could select an incentive compatible relational contract that extracts a bribe with the credible threat to reveal the evidence otherwise. Under such a relational contract, the blackmailer's negotiation payoffs in state x_0 would be positive. Yet, positive negotiation payoffs imply that a contract in which the evidence is revealed (off the equilibrium path) would not be incentive compatible.

5.2 Repeated games

A repeated game with transfers corresponds to the special case that there is just a single state.

Proposition 3. *In a repeated game, negotiation payoffs are regular and their sum is equal to the highest joint PPE payoff of the repeated game given an adjusted discount factor of*

$$\tilde{\delta} = (1 - \rho)\delta.$$

¹⁵The blackmailer could extort larger payments if the game allows to conduct brinkmanship (see e.g. Schelling, 1960 or Schwarz and Sonin, 2007). The Blackmailer needs an observable action that reveals the evidence with positive probability smaller 1. For example, he could leave an envelope with a copy of the evidence addressed to a journalist next to a postal box on the street and then informing the target about it. There is a positive probability that the envelope is still lying on the street if he comes to fetch it, but the envelope might already have been put into the postal box by some helpful minded pedestrian.

This result simply reflects the well known fact that in repeated games, a probability that the relationships ends is equivalent to a lower discount factor: new negotiation essentially constitute a termination and restart of the relationship.

For the case of perfect monitoring and pure strategies, negotiation payoffs of canonical RNE have a particularly simple structure. They split joint payoffs according to a generalized Nash bargaining solution with the threat point given by each player's stage game best reply payoff against the action profile played in her punishment phase.

Proposition 4. *In a repeated game with perfect monitoring expected negotiation payoffs of a canonical RNE with a pure action plan $(a^k)_k$ satisfy*

$$r_i = \pi_i^*(a^i) + \beta_i(\Pi(a^e) - \sum_{j=1}^n \pi_j^*(a^j)) \quad (3)$$

where $\pi^*(a)$ denotes player i 's stage game best-reply payoff

$$\pi_i^*(a) = \max_{\hat{a}_i \in A_i} \pi_i(\hat{a}_i, a_{-i}).$$

5.3 Principal-agent relationship with endogenous inside option

Our main example studies the effects of negotiations and hold-up in a classical relational contracting application: a principal-agent relationship. We first consider a simple repeated game, similar to the ones studied by MacLeod and Malcomson (1998) and Levin (2003). In each period, the agent can choose effort $e \in [0, \bar{e}]$, which determines the value of its service to the principal. The principal's stage game payoff is given by $\max\{x, e\}$. The parameter $x \in (0, \bar{e})$ describes the principal's inside option, i.e. her payoff in a period in which the agent chooses too little effort.¹⁶

The agent has effort costs $k(e)$ that are strictly increasing in e and satisfy $k(e) = 0$, i.e. the agent has an inside option fixed to zero. The joint stage game payoffs $\Pi(e, x) = \max\{x, e\} - k(e)$ shall be strictly increasing in effort for all $e > x$. All transfers and service levels are perfectly observed by the principal and agent, but no enforceable contracts can be written.

As explained in Section 5.2, we can find negotiation payoffs of this repeated game by first solving for optimal simple subgame perfect equilibria given an adjusted discount factor of $\tilde{\delta} = \delta(1 - \rho)$. The lowest punishment payoffs that can be implemented are given by players' inside options. In a simple equilibrium that implements positive effort, the principal pays the agent an bonus if the agent chooses the effort that was agreed upon. Using e.g. the results by Goldlücke and

¹⁶Example 5.4 provides a comparison between inside and outside options.




Figure 3: Pareto frontiers of SPE payoffs in the repeated principal-agent game with discount factor $\tilde{\delta}$ for a low and high inside option x_L and x_H , respectively, with $x_H < \bar{U}(x_H) < \bar{U}(x_L)$.

Kranz (2012) or Proposition 13 in Appendix A, one finds that effort e can be implemented in the equilibrium phase of a simple equilibrium if and only if

$$e - \tilde{\delta}^{-1}k(e) \geq x. \quad (4)$$

The highest effort level that can be implemented is decreasing in the principal's inside option x . That is because under a higher inside option, only weaker punishments can be imposed on a principal who deviates from a bonus payment with the consequence that only lower bonus payments and lower effort levels can be implemented on the equilibrium path.

Figure 5.3 illustrates the corresponding Pareto frontiers of SPE equilibrium payoffs for a low and a high inside option, respectively. $\tilde{U}(x)$ shall denote the highest joint SPE payoff of the repeated game with inside option x given discount factor $\tilde{\delta}$. The principal's (expected) negotiation payoffs are given by

$$\tilde{r}_1(x) = (1 - \beta_1)x + \beta_1\tilde{U}(x) \quad (5)$$

A lower inside option can have two opposing effects on the principal's negotiation payoff. The positive effect is that due to more effective punishments, the size of the joint payoff $\tilde{U}(x)$ can increase. The negative effect is that due to the weaker bargaining position, the principal extracts only a smaller share of the joint payoff $\tilde{U}(x)$. The positive effect dominates the negative effect only if the principal's bargaining weight β_1 is sufficiently large.

Possibility to destroy the inside option Consider now a simple variation with an endogenous inside option. The principal starts with a high inside option x_H but has the possibility to destroy it. Destruction shall reduce the inside option permanently and non-reversibly to a lower level x_L from the next period onwards

and involve no direct costs. In other words, the principal makes herself permanently more vulnerable and more dependent on the agent by destroying her inside option.

Proposition 5. *If a lower inside option is socially more efficient in the repeated game, i.e. $\tilde{U}(x_L) > \tilde{U}(x_H)$, then the principal destroys her inside option in every Pareto-optimal subgame perfect equilibrium (including the RNE for $\rho = 0$).*

Recall that $\tilde{U}(x_L) > \tilde{U}(x_H)$ always holds, except for the uninteresting cases that in a Pareto-optimal equilibrium the agent chooses zero effort or first best effort \bar{e} can be implemented for both x_L and x_H . Destruction of the inside option can be interpreted as a socially beneficial investment which allows to implement a higher joint surplus in the future relationship. Proposition 5 states that Pareto-optimal SPE fully circumvent any hold up problem: the principal always undertakes the socially efficient long-run decision to make herself more vulnerable towards the agent. The absence of a hold-up problem is based on a similar intuition as in Section 2. After the inside option is destroyed, players can flexibly implement any desired continuation payoff from the Pareto-frontier of SPE payoffs of the repeated game with inside option x_L . All that matters for Pareto-optimal SPE is the fact that the continuation payoff set under x_L is strictly larger than under x_H (recall Figure 5.3). Players can permanently, perfectly coordinate away from any exploitation of the principal's increased vulnerability on the equilibrium path.

A positive negotiation probability essentially boils down to the following idea: *After some while, vulnerabilities will be exploited in economic relationships and history-independent bargaining power will be reflected in continuation payoffs.* An agent will exploit a lower inside option of the principal because in new negotiations the agent selects relational contracts that push the principal's payoff down to the inside option. The following result shows that when fixing the adjusted discount factor, the principal's willingness to destroy the inside option is the lower, the higher is the negotiation probability.

Proposition 6. *Holding the adjusted discount factor $\tilde{\delta}$ fixed, the principal destroys her inside option in a RNE if and only if $\tilde{U}(x_H) \leq \tilde{U}(x_L)$ and the negotiation probability is below a critical value $\bar{\rho}(\beta_1)$ that increases in the principal's bargaining weight β_1 .*

Of course, the principal's willingness to destroy the inside option also depends on the trade-off between the resulting increase in the total surplus and the reduction of her share of that surplus in future negotiations. If the principal has a higher bargaining weight β_1 , she can generally appropriate a larger share of the total surplus and her inside option becomes less relevant for the negotiation outcome; the principal is then more willing to destroy the inside option.

The trade-off between the increase in joint surplus and weakening of the principal's bargaining position takes a particularly intuitive form in the limit in which all adjusted discounting is only due to the negotiation probability.

Proposition 7. *Fix a negotiation probability $\rho \in (0, 1)$ and consider the limit $\delta \rightarrow 1$. The principal destroys her inside option if and only if her negotiation payoff in the repeated game with fixed inside option is larger under the low inside option, i.e. $\tilde{r}_1(x_H) \leq \tilde{r}_1(x_L)$.*

With a positive negotiation probability continuation, payoffs will after some while be given by the negotiation payoffs that reflect in each state player's history independent bargaining power. In the limit case of $\delta \rightarrow 1$ payoffs in intermediate periods get zero weight and only the resulting negotiation payoffs will be relevant for the principal's choice of state. The principal thus solves exactly the same trade-off between social benefits and weakening of her bargaining position that determines whether she would be better off under a low inside option in the repeated game with exogenously given inside options. We get a very simple formula for this trade-off if bargaining weights are equal, i.e. $\beta_1 = \beta_2 = \frac{1}{2}$. The principal then destroys the inside option whenever

$$\tilde{U}(x_L) - \tilde{U}(x_H) \geq x_H - x_L.$$

The increase in joint payoffs must exceed the decrease in the principal's inside option.

As a final observation for this example, we note that the game has a unique Markov Perfect equilibrium, which is also the unique RNE for $\rho = 1$. This limit case does not illuminate any interesting trade-offs. The agent will simply never choose positive efforts and the principal will consequently never destroy her inside option.

5.4 Inside options versus outside options

An important insight of non-cooperative bargaining models with enforceable contracts is the distinction between inside options, which describe the payoffs during periods of disagreement within the relationship, and outside options, which describe the payoffs if the relationship breaks up without forming a contract. The famous outside option principle, states that outside options should only influence bargaining outcomes if they are binding while otherwise only inside options are relevant, see e.g. Binmore et. al. (1989).

The difference between outside options and inside options can have important implications for hold-up problems, see e.g. De Meza and Lockwood (1998). In contrast, in the relational contracting literature these differences are mostly irrelevant and to the best of our knowledge, an outside option principle for relational contracts has not yet been established or studied.

This example illustrates how repeated negotiations naturally extend the outside option principle to relational contracting. Consider a variation of the repeated

principal agent game from the previous section. Within the relationship the principal's and agent's stage game payoffs shall be given by

$$\begin{aligned}\pi_1 &= \pi_1^{io} + e, \\ \pi_2 &= \pi_2^{io} - k(e).\end{aligned}$$

As before, $e \in [0, \bar{e}]$ denotes the agent's effort and $k(e)$ an increasing cost function with $k(0) = 0$. The payoff vector π^{io} denotes players' inside options and describes the payoffs in case zero quality is delivered.

In each period the principal and agent can also decide to break-up their relationship. If both want to break up their relationship, the break-up is permanent and player i gets in the current and all future periods an outside option payoff of π_i^{oo} . We assume that $\pi_i^{io} \leq \pi_i^{oo}$, i.e. both players prefer a break-up compared to staying in the relationship while never trading with each other. To rule out that a player is indifferent between breaking up or not if the other player wants to break-up, we assume that if just one player wants to break up, there is a very small positive probability $\varepsilon \approx 0$ that the break up is not successful and players remain in the relationship next period.

Proposition 8. *Let \bar{U} denote the joint payoffs of the RNE. If the negotiation probability is zero, expected negotiation payoffs are only determined by the outside options and given by*

$$r_i^{oo} = \pi_i^{oo} + \beta_i(\bar{U} - \pi_1^{oo} - \pi_2^{oo}). \quad (6)$$

Given a fixed positive negotiation probability $\rho > 0$ and the limit $\delta \rightarrow 1$, expected negotiation payoffs satisfy instead the outside option principle: unless the outside options are binding, they are solely determined by the inside options and are given by

$$r_i^{io} = \pi_i^{io} + \beta_i(\bar{U} - \pi_1^{io} - \pi_2^{io}). \quad (7)$$

These results are straightforward. The lowest subgame perfect equilibrium payoff of each player is given by her outside option. In the absence of repeated negotiation, the expected negotiation payoffs r_i^{oo} therefore simply split the joint surplus as in the Nash bargaining solution with the outside options as threat point.

The payoffs r_i^{io} correspond to the expected negotiation payoffs inside the relationship, i.e. in a game in which for no player it is possible or credible to break up the relationship. In the case $\rho > 0$ and $\delta \rightarrow 1$, continuation payoffs are always approximately equal to subsequent negotiation payoffs. Hence, breaking up the relationship would only be incentive compatible if at least for one player i the outside option payoff is larger than her negotiation payoff inside the relationship r_i^{io} . Otherwise, outside options have no influence on the equilibrium outcome, i.e. the outside option principle holds.

5.5 A simple arms race

The final example illustrates how repeated negotiations can lead to excessive investments into means to harm other player in order to gain bargaining power. It also illustrates that equilibrium payoffs do not necessarily decrease in the negotiation probability but may also increase in it.

Consider two players, e.g. different countries. In certain periods, one or both players can have the opportunity to spend an amount of money $b > 0$ in order to try to acquire a weapon. Whether that opportunity arises or an attempted acquisitions is successful may be determined in a stochastic fashion. Once a weapon has been successfully acquired, players can use it in later periods at some cost $c > 0$ to inflict a damage $d > 0$ on the other player. There are no direct benefits from using a weapon. Players can acquire at most one weapon and the weapon can be used as often as desired. If no weapons are bought or used, players get zero payoffs.

Proposition 9. *In the unique Markov perfect equilibrium outcome, as well as, in all Pareto-optimal SPE outcomes no weapons are bought or used; this also holds true for the corresponding RNE given negotiation in every period ($\rho = 1$) or no repeated negotiation ($\rho = 0$). In contrast, for intermediate negotiation probabilities, it can be the case that there is a unique RNE outcome in which one or both players spend money to acquire weapons.*

The fact that weapons are not build in the two extreme cases of no repeated negotiation or negotiation in every period has different reasons.

That costly acquisition of weapons on the equilibrium path cannot be part of a Pareto efficient SPE, is evident, since joint payoffs are maximized if no investment cost are incurred. Non-investment can e.g. be sustained if players coordinate to ignore forever any threat to use weapons.

If negotiations occur in every period (corresponding to the unique MPE), weapons will never be used since usage is costly and has by definition no impact on the future state, i.e. attacks cannot induce any future payments.

Under intermediate negotiation probabilities the two factors that block weapon acquisition can be relaxed simultaneously. Incentive compatible relational contracts in which weapons are used until the other player makes an appeasement payment exist as long as the negotiation probability is not too high. If at the same time the negotiation probability is sufficiently high, players may not be able to prevent themselves from acquiring weapons, since for each player the resulting gain in future negotiation payoffs is simply too attractive.

The example illustrates that to prevent an arms race, unfortunately, it is not sufficient that there are no direct gains from using weapons. Struggles for power are simply a natural phenomenon under positive negotiation probabilities. The induced socially inefficient investments can be interpreted as a particular incarnation of a hold-up problem.

6 Extensions

6.1 State-dependent Negotiation Probabilities

A straightforward extension of our basic model is to allow for heterogeneous negotiation probabilities that can depend on the current state. In the proofs in Appendix A, we directly consider this more general model and find

Proposition 10. All results in Section 4, including the existence theorem for canonical RNE, extend to the case that negotiation probabilities depend on the state x .

State-dependent negotiation probabilities allow for a rich variety of variants of our basic model that automatically satisfy the key existence and characterization results.

One example is to model incompleteness of relational contracts with respect to certain contingencies. To model the fact that some state \tilde{x} is not considered in an initial relational contract, one can assign a negotiation probability of 1 to that state while having smaller or zero negotiation probabilities for states that have been initially considered.

Another possibility is a model in which negotiation probabilities are reduced over time as players interact more often. This can reflect the transformation of relational contracts from initially very loose informal agreements into well established social norms. This is a special case of our general framework since one can simply augment the state space by an additional dimension that keeps track of the current negotiation probability.

6.2 Endogenous Negotiations

By modifying the original stochastic game, state-dependent negotiation probabilities also allow variants of our model in which new negotiations only take place if some player actively attempts to force negotiations.

Consider the following modified game. The action space in all states is augmented by a set of binary choices “do or do not attempt new negotiations if in the next period state x is reached” for each player. Furthermore, every state $x \in X$ of the original game is split into two states: x_0 and x_1 . In state x_0 the negotiation probability is 0 and in state x_1 it is 1. The two states have identical payoff function, action space and transition probabilities. We call x the base state and x_0 and x_1 its variants. State transitions are as follows. First the new base state x is determined using the same transition probabilities as in the original game. If no player has attempted negotiation for state x then state x_0 is reached, if all players have attempted negotiation then state x_1 is reached and if some but not all players have attempted new negotiations, state x_1 is reached with probability

ρ and otherwise state x_0 is reached. Hence, in the game with endogenous negotiation, ρ measures the probability that a player is able to force new negotiations against the will of other players.

We say a canonical contract profile σ in the game with exogenous negotiation is behavioral equivalent to a canonical contract profile $\tilde{\sigma}$ in a corresponding game with endogenous negotiation if for every base state x , the variant x_1 is reached with probability ρ and the following values are the same as in the state x of the original game: the action profiles and transfers for the punishment and equilibrium phases in both state variants and the negotiation payoffs in state x_1 .

Proposition 11. *Assume unsuccessful negotiation attempts are unobservable. For every incentive compatible canonical contract profile without money burning in the original game with exogenous negotiation there is a behavioral equivalent incentive compatible canonical contract profile in the game with endogenous negotiation attempts.*

While Proposition 11 illustrates common aspects, the outcomes in the models with exogenous and endogenous negotiation are not in general equivalent. For example, in a game with imperfect public monitoring, it can become optimal under an exogenous negotiation probability to burn money in the equilibrium phase as a collective punishment for bad outcomes. Yet, with endogenous negotiation all players may then want to newly negotiate the contract.

The special case that new negotiations only take place if all players attempt negotiations, i.e. $\rho = 0$, can be considered as an extension of renegotiation-proofness concepts from repeated games to stochastic games with transfers and imperfect monitoring. We are not aware of any previous formalization, characterization or existence result of renegotiation-proofness for stochastic games.

If one assumes that unsuccessful negotiation attempts are publicly observed and can thus be punished, it is easier to deter negotiations on the equilibrium path. Still, the option to attempt negotiations will typically increase the lowest punishment payoff that can be imposed on a player. In a similar fashion as in our baseline model, the opportunity to attempt new negotiations reduces the flexibility of relational contracting and can cause hold-up problems.

Bibliography

Asheim, G. B. 1991. “Extending Renegotiation-proofness to Infinite Horizon Games.” *Games and Economic Behavior* 3 (3): 278–294.

Avery, C., and P. B. Zemsky. 1994. “Money Burning and Multiple Equilibria in Bargaining.” *Games and Economic Behavior* 7 (2): 154–68.

Baker, G., R. Gibbons, and K.J. Murphy. 2002. “Relational Contracts and the Theory of the Firm.” *The Quarterly Journal of Economics* 117 (1): 39.

- Baliga, S., and R. Evans. 2000. "Renegotiation in Repeated Games with Side-payments." *Games and Economic Behavior* 33 (2): 159–176.
- Binmore, K., A. Shaked, and J. Sutton. 1989. "An Outside Option Experiment." *The Quarterly Journal of Economics* 104 (4): 753–770.
- Blonski, M. and G. Spagnolo, 2007. "Relational Contracts and Property Rights." Working Paper.
- Chatterjee, K., and L. Samuelson. 1990. "Perfect Equilibria in Simultaneous-offers Bargaining." *International Journal of Game Theory* 19 (3): 237–267.
- Che, Y. K., and D. B Hausch. 1999. "Cooperative Investments and the Value of Contracting." *American Economic Review* 89 (1): 125–147.
- Che, Y. K., and J. Sákovics. 2004. "A Dynamic Theory of Holdup." *Econometrica* 72 (4): 1063–1103.
- Ching, A.T. 2010. "A Dynamic Oligopoly Structural Model for the Prescription Drug Market After Patent Expiration." *International Economic Review* 51 (4): 1175–1207.
- Doornik, K. 2006. "Relational Contracting in Partnerships." *Journal of Economics & Management Strategy* 15 (2): 517–548.
- Douglas Bernheim, B., and D. Ray. 1989. "Collective Dynamic Consistency in Repeated Games." *Games and Economic Behavior* 1 (4): 295–326.
- Dutta, P. K. 1995. "A Folk Theorem for Stochastic Games." *Journal of Economic Theory* 66 (1): 1–32.
- Edlin, A. S., and S. J. Reichelstein. 1996. "Holdups, Standard Breach Remedies, and Optimal Investment." *American Economic Review* 86 (3): 478–501.
- Ellingsen, T., and M. Johannesson. 2004. "Is There a Hold-up Problem?" *The Scandinavian Journal of Economics* 106 (3): 475–494.
- Ellingsen, T., and J. Robles. 2002. "Does Evolution Solve the Hold-up Problem?" *Games and Economic Behavior* 39 (1): 28–53.
- Farrell, J., and E. Maskin. 1989. "Renegotiation in Repeated Games." *Games and Economic Behavior* 1 (4): 327–360.
- Fudenberg, D., D. Levine, and E. Maskin. 1994. "The Folk Theorem with Imperfect Public Information." *Econometrica: Journal of the Econometric Society*: 997–1039.
- Fudenberg, D., and Y. Yamamoto. 2011. "The Folk Theorem for Irreducible Stochastic Games with Imperfect Public Monitoring." *Journal of Economic Theory*.
- Goldluecke, S., and S. Kranz. 2012. "Infinitely Repeated Games with Public Monitoring and Monetary Transfers." *Journal of Economic Theory* 147 (3): 1191–1221.
- Grout, P. A. 1984. "Investment and Wages in the Absence of Binding Contracts: A Nash Bargaining Approach." *Econometrica: Journal of the Econometric Society*: 449–460.
- Halonen, M. 2002. "Reputation and the Allocation of Ownership." *The Economic Journal* 112 (481): 539–558.

- Harrington, J.E., and A. Skrzypacz. 2011. "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices." *The American Economic Review* 101 (6): 2425–2449.
- Harrington Jr, J.E., and A. Skrzypacz. 2007. "Collusion Under Monitoring of Sales." *The RAND Journal of Economics* 38 (2): 314–331.
- Hart, O., and J. Moore. 1988. "Incomplete Contracts and Renegotiation." *Econometrica: Journal of the Econometric Society* 56 (4): 755–785.
- Hörner, J., T. Sugaya, S. Takahashi, and N. Vieille. 2011. "Recursive Methods in Discounted Stochastic Games: An Algorithm for $\delta \rightarrow 1$ and a Folk Theorem." *Econometrica* 79 (4): 1277–1318.
- Klimenko, M., G. Ramey, and J. Watson. 2008. "Recurrent Trade Agreements and the Value of External Enforcement." *Journal of International Economics* 74 (2): 475–499.
- Kranz, S. 2012. "Discounted Stochastic Games with Voluntary Transfers." Cowles Foundation Discussion Paper No. 1847.
- Kranz, S., and S. Ohlendorf. 2009. "Renegotiation-proof Relational Contracts with Side Payments." SFB/TR 15 Discussion Paper 259.
- Levin, J. 2002. "Multilateral Contracting and the Employment Relationship." *The Quarterly Journal of Economics* 117 (3): 1075.
- Macaulay, S. 1963. "Non-contractual Relations in Business: A Preliminary Study." *American Sociological Review*: 55–67.
- MacLeod, W. B., and J. M. Malcomson. 1989. "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment." *Econometrica: Journal of the Econometric Society* 57 (2): 447–480.
- Malcomson, J. M. 2010. "Relational Incentive Contracts." Mimeo.
- Maskin, E., and J. Tirole. 2001. "Markov Perfect Equilibrium:: I. Observable Actions." *Journal of Economic Theory* 100 (2): 191–219.
- De Meza, D., and B. Lockwood. 1998. "Does Asset Ownership Always Motivate Managers? Outside Options and the Property Rights Theory of the Firm." *The Quarterly Journal of Economics* 113 (2): 361.
- Miller, D.A., J. Watson, 2011. "A Theory of Disagreement in Repeated Games with Bargaining.", mimeo.
- Nöldeke, G., and K. M Schmidt. 1995. "Option Contracts and Renegotiation: a Solution to the Hold-up Problem." *The RAND Journal of Economics* 26 (2): 163–179.
- Ohlendorf, Susanne. 2009. "Expectation Damages, Divisible Contracts, and Bilateral Investment." *American Economic Review* 99 (4) (August): 1608–1618. doi:10.1257/aer.99.4.1608.
- Pakes, A., and P. McGuire. 1994. "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model." *The Rand Journal of Economics*: 555–589.
- . 2001. "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'curse' of Dimensionality." *Econometrica* 69 (5): 1261–1281.

Pitchford, R., and C. M. Snyder. 2004. “A Solution to the Hold-up Problem Involving Gradual Investment.” *Journal of Economic Theory* 114 (1): 88–103.

Ramey, G., and J. Watson. “Contractual Fragility, Job Destruction, and Business Cycles.” *The Quarterly Journal of Economics* 112, no. 3 (1997): 873–911.

Rayo, L. 2007. “Relational Incentives and Moral Hazard in Teams.” *The Review of Economic Studies* 74 (3): 937.

Rubinstein, A. 1982. “Perfect Equilibrium in a Bargaining Model.” *Econometrica: Journal of the Econometric Society*: 97–109.

Schelling, T. C. 1960. “The Strategy of Conflict.” Cambridge, MA: Harvard Univ.

Schmidt, K. M., and M. Schnitzer. 1995. “The Interaction of Explicit and Implicit Contracts.” *Economics Letters* 48 (2): 193–199.

Schwarz, M., and K. Sonin. 2008. “A Theory of Brinkmanship, Conflicts, and Commitments.” *Journal of Law, Economics, and Organization* 24 (1): 163–183.

Sobel, M. J. 1971. “Noncooperative stochastic games.” *The Annals of Mathematical Statistics*: 1930–1935.

Tröger, T. 2002. “Why Sunk Costs Matter for Bargaining Outcomes: An Evolutionary Approach.” *Journal of Economic Theory* 102 (2): 375–402.

Appendix A: Existence and Characterization of canonical RNE

Proofs for results in Section 4 and 6

This appendix proves the results in Section 4 and 6 and develops further results for characterizing and understanding the structure of RNE. Recall that $\mathbf{g}(\cdot|\boldsymbol{\sigma}, \mathbf{r})$ the payoffs of the relational contracts in the corresponding truncated games and $G^\sigma : \mathcal{R} \rightarrow \mathcal{R}$ is the following functional operator:

$$G^\sigma(\mathbf{r}) \equiv \mathbf{g}(\cdot|\mathbf{r}, \boldsymbol{\sigma}).$$

Lemma 2. *G^σ is monotone increasing and in the metric space (\mathcal{R}, d_∞) a contraction mapping that has a unique fixed point given by the negotiation payoffs induced by $\boldsymbol{\sigma}$ in the original game. Therefore*

$$\mathbf{r} = \mathbf{g}(\cdot|\mathbf{r}, \boldsymbol{\sigma}) \Leftrightarrow \mathbf{r} = \mathbf{r}(\cdot|\boldsymbol{\sigma}).$$

Proof. First note that $\mathbf{r}(\cdot|\boldsymbol{\sigma}) = \mathbf{g}(\cdot|\mathbf{r}(\cdot|\boldsymbol{\sigma}), \boldsymbol{\sigma})$ follows straightforward from our definitions. If G^σ is a contraction mapping, the contraction mapping theorem implies that $\mathbf{r}(\cdot|\boldsymbol{\sigma})$ is the unique fixed point of G^σ . We show that G^σ is a contraction mapping using Blackwell’s sufficient conditions. The first condition is that

G^σ is monotone increasing in the following form: if two negotiation payoffs \mathbf{r} and $\tilde{\mathbf{r}}$ satisfy $\mathbf{r}_j^i(x) \leq \tilde{\mathbf{r}}_j^i(x) \forall i, j, x$ then $\mathbf{g}_j^i(x|\mathbf{r}, \sigma) \leq \mathbf{g}_j^i(x|\tilde{\mathbf{r}}, \sigma) \forall i, j, x$. Verbally, payoffs of the truncated game are increasing in the negotiation payoffs if the strategy profile is hold fixed. This monotonicity condition is obviously satisfied.

The discounting condition requires that there exist a scalar $\gamma \in (0, 1)$ such that for any constant $K \geq 0$ and all i, j, x

$$\mathbf{g}_j^i(x|\hat{\mathbf{r}} + K, \sigma) \leq \mathbf{g}_j^i(x|\hat{\mathbf{r}}, \sigma) + \gamma K$$

Average discounted payoffs of the truncated game can increase at most by δK if negotiation payoffs increase by K . because transition to an absorbing state can only occur after period 1. The discounting condition is therefore satisfied with $\gamma = \delta$. \square

For convenience, we state the following direct consequence of Lemma 2.

Lemma 3. *Two contract profiles σ and $\tilde{\sigma}$ induce the same negotiation payoffs, i.e.*

$$\mathbf{r}(\cdot|\sigma) = \mathbf{r}(\cdot|\tilde{\sigma}),$$

if and only if they implement the same payoffs in the truncated games given $\mathbf{r}(\cdot|\sigma)$, i.e.

$$\mathbf{g}(\cdot|\mathbf{r}(\cdot|\sigma), \sigma) = \mathbf{g}(\cdot|\mathbf{r}(\cdot|\sigma), \tilde{\sigma})$$

Proposition 2. *For any incentive compatible contract profile σ there exists an incentive compatible canonical contract profile $\tilde{\sigma}$ which has the same negotiation payoffs $\mathbf{r}(\cdot|\sigma) = \mathbf{r}(\cdot|\tilde{\sigma})$.*

Proof. It follows from Theorem 1 that for every state x there exists an optimal simple equilibrium $\tilde{\sigma}^{(i,x)}$ that by setting appropriate upfront transfers implement the same payoffs as $\sigma^{(i,x)}$ in the truncated game $\Gamma(\mathbf{r}(\cdot|\sigma), x)$, A profile $\tilde{\sigma}$ of such optimal simple renegotiation outcomes thus satisfies

$$\mathbf{g}^i(x|\mathbf{r}(\cdot|\sigma), \sigma) = \mathbf{g}^i(x|\mathbf{r}(\cdot|\sigma), \tilde{\sigma}) \forall i = 1, \dots, n, x \in X$$

It follows from Lemma 3 that

$$\mathbf{r}(\cdot|\sigma) = \mathbf{r}(\cdot|\tilde{\sigma}).$$

and $\tilde{\sigma}$ is therefore indeed incentive compatible. \square

The following result establishes that player i always has an incentive compatible simple relational contract in state x if all other offered relational contracts are simple strategy profiles. The proof specifies a modified stochastic game and then exploits Sobel's (1971) existence result of Markov perfect equilibria in stochastic games with finite action and state spaces.

Lemma 4. *Assume the pure action space is finite. For every profile of simple relational contracts $\sigma_{-(i,x)}$, there always exists an incentive compatible simple relational contract $\sigma_{(i,x)}$ for player i in state x .*

Proof. We show that there is always an incentive compatible relational contract $\sigma_{(i,x)}$ that forms a MPE of the corresponding truncated game. A MPE is just a special form of a simple equilibrium in which players conduct no payments and equilibrium phase and punishment phase policies coincide.

For any pair (i, x) of player and state and given simple contract profile $\sigma_{-(i,x)}$, consider the following modified stochastic game $G(\sigma_{-(i,x)})$. The modified states are described by (x', i_n, x_n) where x' corresponds to the state of the original game, x_n is the state in which the previous negotiation took place and i_n describes the player that has selected the current relational contract. In modified states (x', i_n, x_n) with $(i_n, x_n) \neq (i, x)$ players have no choice of actions: play automatically proceeds as in the equilibrium phase of the relational contract $\sigma_{(i_n, x_n)}$ in the original game with the same payoffs. In modified states (x', i_n, x_n) with $(i_n, x_n) = (i, x)$, the players' action space is the same as the action space of the original game in state x' , yet no transfers are possible, i.e. the action space of the modified game is finite. Stage game payoff functions in a modified state (x', i_n, x_n) are as the payoff functions in state x' in the original game. State transitions of the component x' of the modified state are like the state transitions in the original game and i_n and x_n are updated when new negotiations take place.

The modified game $G(\sigma_{-(i,x)})$ is a stochastic game with a finite action space and a finite state space. Sobel (1971) has shown that it must have a MPE. It follows straightforward from the construction that the MPE of $G(\sigma_{-(i,x)})$ constitutes an incentive compatible relational contract for player i in state x given $\sigma_{-(i,x)}$: no player has an incentive to deviate, holding future negotiation outcomes fixed. \square

The following result makes use of the fact that the set of simple equilibria is compact.

Lemma. *If for a given contract profile $\sigma_{-(i,x)}$, player i has an incentive compatible relational contract in state x then she has an incentive compatible simple contract that gives her the supremum of her continuation payoffs across all incentive compatible relational contracts.*

Proof. We denote by $\Sigma^*(\sigma_{-(i,x)})$ the set of incentive compatible relational contracts for player i in state x given $\sigma_{-(i,x)}$. Let

$$\bar{r}_i^i = \sup_{\tilde{\sigma}_{(i,x)} \in \Sigma_i^*(\sigma_{-(i,x)})} r_i^i(x | \tilde{\sigma}_{(i,x)}, \sigma_{-(i,x)})$$

denote the supremum of player i 's payoffs in state x that can be implemented with incentive compatible relational contracts. Using similar arguments as in the proof of Lemma 2, one can show that every incentive compatible relational

contract chosen by player i in state x can be replaced by an incentive compatible canonical relational contract that has the same negotiation payoffs.

Let $\{\tilde{\sigma}_{(i,x)}(m)\}_{m=1}^{\infty}$ and $\{\tilde{r}_i^i(m)\}_{m=1}^{\infty}$ be sequences of incentive compatible simple optimal relational contracts and corresponding payoffs for player i such that $\{\tilde{r}_i^i(m)\}_{m=1}^{\infty}$ converges towards \bar{r}_i^i . Since $\{\tilde{\sigma}_{(i,x)}(m)\}_{m=1}^{\infty}$ is a sequence in a compact space, it must have a converging subsequence and we denote its limit by $\bar{\sigma}_{(i,x)}$. It is straightforward that $\bar{\sigma}_{(i,x)}$ must be an optimal simple equilibrium that implements \bar{r}_i^i in the truncated game $\Gamma(\mathbf{r}(\bar{\sigma}_{(i,x)}, \boldsymbol{\sigma}_{-(i,x)}), x)$. That is because payoffs are continuous in mixing probabilities over actions and in payments and all incentive constraints of a simple equilibrium consists of weak inequalities (see Kranz, 2012, for a detailed description of those incentive constraints) \square

We now prove the main existence result. The main idea of the proof is to construct a class of auxiliary games and to connect existence of negotiation equilibria with existence of Nash equilibria in those auxiliary games.

Theorem 2. *If the action space is finite and mixed actions are allowed, a canonical RNE exists.*

Proof. Before proving the result, we prove a simpler result that establishes that there exists a RNE in which all relational contracts are simple contracts. This is not yet a canonical RNE, since we do not yet require that all relational contracts are optimal simple contracts of the truncated game that only differ in their upfront transfers.

Existence of a RNE in simple contracts We first introduce some definitions for the original stochastic game. A point of play $\mathfrak{h}_{(i,x)} = (s, k, y, x', \tilde{x}, i, x)$ shall be a vector that contains all aspects of a history in the original stochastic game that is relevant for players' continuation payoffs given a simple contract profile. A point of play consists of the current stage s (transfer or action stage), the phase $k \in \{e, 1, \dots, n\}$, the previously realized signal y , the current state x' , the previous state \tilde{x} , and a pair of player and state (i, x) that identifies the current relational contract. $\mathfrak{H}((i, x))$ shall denote the finite set of all points of play in which the current relational contract has been selected by player i in state x . For a strategy profile described by a simple contract profile $\boldsymbol{\sigma}$, let

$$\mathbf{u}_j(\mathfrak{h}, \boldsymbol{\sigma})$$

denote player j 's expected continuation payoff at a point of play \mathfrak{h} . For $\mathfrak{h} \in \mathfrak{H}((i, x))$ let

$$\mathbf{u}_j^*(\mathfrak{h}, \boldsymbol{\sigma})$$

denote the maximum of player j 's continuation payoffs if she is allowed to deviate from the relational contract $\boldsymbol{\sigma}_{(i,x)}$ but is not allowed to deviate from the other relational contracts. With some abuse of notation, we call $\mathbf{u}_j^*(\mathfrak{h}, \boldsymbol{\sigma})$ player j 's best reply payoff. Payoffs $\mathbf{u}_j(\mathfrak{h}, \boldsymbol{\sigma})$ are linear in $\boldsymbol{\sigma}$ and best-reply payoffs $\mathbf{u}_j^*(\mathfrak{h}, \boldsymbol{\sigma})$ are

convex in σ . Correspondingly, the gains from an optimal deviation $u_j^*(h, \sigma) - u_j(h, \sigma)$ are convex in σ . We denote by

$$\Delta_{(i,x)}(\sigma) = \max_{h \in \mathfrak{H}(i,x)} \max_{j \in \{1, \dots, n\}} (u_j^*(h, \sigma) - u_j(h, \sigma))$$

the highest gain from an optimal deviation from the relational contract $\sigma_{(i,x)}$ across all players and corresponding points of play. As a maximum over a finite number of continuous and convex functions, $\Delta_{(i,x)}(\sigma)$ is quasi-convex and continuous.

We now specify a simultaneous move auxiliary game. Each combination (i, x) of a player and state of the original stochastic game constitutes a player of the auxiliary game. The compact action space of an auxiliary player (i, x) is given by the set of simple contracts of the truncated game starting in state x . Hence, a simple contract profile σ specifies a strategy profile of this simultaneous move auxiliary game. A particular auxiliary game is described by a parameter $\varepsilon > 0$. Payoffs of auxiliary player (i, x) are given by

$$\pi_{(i,x)}(\sigma) = \min\{r_i^i(x|\sigma), K - \frac{1}{\varepsilon} \Delta_{(i,x)}(\sigma)\}$$

where K is positive constant that is larger than the maximum (average discounted) continuation payoff a player could achieve in the original game. As long as the incentive constraints of the relational contract $\sigma_{(i,x)}$ are not violated by too large an amount, i.e. $\Delta_{(i,x)}(\sigma)$ is not too large, the auxiliary player's payoff $\pi_{(i,x)}(\sigma)$ is given by player i 's continuation payoff when selecting $\sigma_{(i,x)}$ in state x in the original game, i.e. by $r_i^i(x|\sigma)$. The smaller is ε , the more severe do violations of the incentive constraints enter the the auxiliary players' payoffs. The negotiation payoffs $r_i^i(x|\sigma)$ are a linear function of σ . As the minimum of a linear and a quasi-concave function, $\pi_{(i,x)}(\sigma)$ is itself a quasi-concave function and also continuous. Since payoffs are quasi-concave and continuous and the action space is compact, we can apply the standard Nash equilibrium existence proof to conclude that the auxiliary game has a Nash equilibrium for all $\varepsilon > 0$.

Consider an infinite sequence of $\{\varepsilon^m\}_{m=1}^{\infty}$ that converges to 0 and a corresponding sequence of auxiliary games with $\varepsilon = \varepsilon^m$. Let $\{\sigma^m\}_{m=1}^{\infty}$ be the corresponding sequence of Nash equilibria of those auxiliary games. As an infinite sequence in a compact space, $\{\sigma^m\}_{m=1}^{\infty}$ has a convergent subsequence. Without loss of generality, we assume that $\{\sigma^m\}_{m=1}^{\infty}$ is already that convergent subsequence and denote its limit by σ^* . It follows from Lemma 4 that every auxiliary player has a simple relational contract $\sigma_{(i,x)}$ that is incentive compatible given $\sigma_{-(i,x)}^*$. Since violations of incentive constraints are exceedingly costly as $\varepsilon \rightarrow 0$, it thus follows from our construction that $\sigma_{(i,x)}^*$ must be incentive compatible given $\sigma_{-(i,x)}^*$ and that σ^* constitutes a RNE.

Existence of canonical RNE To prove existence of canonical RNE, we use a variation of the auxiliary game developed above. In addition to the auxiliary

players (i, x) there will be a *canonical player* indexed by c . The canonical player specifies a contract profile $\tilde{\sigma}$ consisting of simple relational contracts that only differ in their upfront payments. As before, each auxiliary player (i, x) chooses a relational contract $\sigma_{(i,x)}$. An auxiliary player's payoffs of the new game shall be given by

$$\pi_{(i,x)}(\sigma, \tilde{\sigma}) = \min\{r_i^i(x|\sigma_{(i,x)}, \tilde{\sigma}_{-(i,x)}), K - \frac{1}{\varepsilon}\Delta_{(i,x)}((\sigma_{(i,x)}, \tilde{\sigma}_{-(i,x)}))\}.$$

The payoff function differs from our previous specification in so far that payoffs and violations of incentive constraints are computed under the assumption that other relational contracts are given by the canonical player's choice $\tilde{\sigma}_{-(i,x)}$ instead of the profile $\sigma_{-(i,x)}$ of the other auxiliary players. $\pi_{(i,x)}(\sigma, \tilde{\sigma})$ is quasi-concave and continuous in both σ and $\tilde{\sigma}$.

The specification of the canonical player's payoff function is slightly more involved. Let $U^{(i,x)}(\tilde{\sigma})$ denote the joint equilibrium phase payoffs of the simple contract $\tilde{\sigma}_{(i,x)}$ in the truncated game $\hat{\Gamma}(r(\cdot|\tilde{\sigma}), x)$. In a similar fashion, we denote by $v_j^{(i,x)}(\tilde{\sigma})$ the corresponding punishment payoffs of player j . Let $V^{(i,x)}(\tilde{\sigma})$ denote the sum of those punishment payoffs across all player j .

The canonical player's payoffs shall be given by

$$\begin{aligned} \pi_{(c)}(\sigma, \tilde{\sigma}) = \min\{ & \sum_{\forall(i,x)} (U^{(i,x)}(\tilde{\sigma}) - V^{(i,x)}(\tilde{\sigma})), \\ & K - \frac{1}{\varepsilon} \max_{\forall j, \forall(i,x)} (r_j^i(x|\sigma_{(i,x)}, \tilde{\sigma}_{-(i,x)}) - r_j^i(x|\tilde{\sigma}))^2, \\ & K - \frac{1}{\varepsilon} \max_{\forall(i,x)} \Delta_{(i,x)}(\tilde{\sigma})\} \end{aligned}$$

Intuitively, the canonical player chooses her contract profile $\tilde{\sigma}$ such that the common structure of actions and non-upfront transfers maximize the sum of gaps between joint equilibrium phase and punishment payoffs in the corresponding truncated games under the restrictions that i) corresponding negotiation payoffs $r_j^i(x|\tilde{\sigma})$ are sufficiently close to the negotiation payoffs $r_j^i(x|\sigma_{(i,x)}, \tilde{\sigma}_{-(i,x)})$ for every player j and each auxiliary player (i, x) , and ii) the relational contracts don't violate the incentive constraints by too large an amount. As ε decreases, the two restrictions become exceedingly tighter.

All payoff functions of the new auxiliary game are continuous and quasi-concave in σ and $\tilde{\sigma}$, which implies that for every $\varepsilon > 0$ the auxiliary game has a Nash equilibrium. Similar as above, let $\{(\tilde{\sigma}^m, \sigma^m)\}_{m=1}^{\infty}$ denote a converging subsequence of such Nash equilibria as $\varepsilon \rightarrow 0$ and let $(\tilde{\sigma}^*, \sigma^*)$ denote its limit point.

It follows again from Lemma4 that in this limit the auxiliary contracts $\sigma_{(i,x)}^*$ are incentive compatible given $\tilde{\sigma}_{-(i,x)}^*$ and maximize the auxiliary player (i, x) 's payoffs across all incentive compatible contracts. Let \hat{r} denote the negotiation

payoffs defined by

$$\hat{r}^i = r^i(\cdot | \sigma_{(i,x)}^*, \tilde{\sigma}_{-(i,x)}^*).$$

It follows from Theorem 1, that there exists a simple contract profile $\hat{\sigma}$ in which the relational contracts only differ in their upfront payments and for each (i, x) the relational contract $\hat{\sigma}_{(i,x)}$ is an optimal simple equilibrium of the truncated game $\hat{\Gamma}(\hat{r}, x)$ that implements payoffs \hat{r}^i . It follows from Lemma 3 that $\hat{\sigma}$ has negotiation payoffs \hat{r} ; therefore it is also an incentive compatible canonical contract profile.

Note that the part

$$\sum_{\forall(i,x)} (U^{(i,x)}(\tilde{\sigma}) - V^{(i,x)}(\tilde{\sigma}))$$

in the canonical player's profit function makes the canonical player want to choose optimal simple contracts of the truncated game $\hat{\Gamma}(\hat{r}, x)$. It therefore follows that also $\tilde{\sigma}^*$ must be an incentive compatible canonical contract profile that implements negotiation payoffs \hat{r} , since otherwise the canonical player would (in the limit) strictly prefer choosing $\hat{\sigma}$ instead of $\tilde{\sigma}^*$. (More precisely, for sufficiently small $\bar{\varepsilon}$ there exist profiles nearby $\hat{\sigma}$ that the canonical player would strictly prefer over $\tilde{\sigma}^m$ for all m such that $\varepsilon^m < \bar{\varepsilon}$). The part,

$$\sum_{\forall(i,x)} (U^{(i,x)}(\tilde{\sigma}) - V^{(i,x)}(\tilde{\sigma}))$$

Since, given $\tilde{\sigma}_{-(i,x)}^*$, the auxiliary player (i, x) has no incentive compatible contract that implements a higher payoff than $\hat{r}_i^i(x)$, $\tilde{\sigma}^*$ constitutes a canonical RNE. \square

The proofs above already imply Proposition 10 in Section 6. We now prove Proposition 11.

Proposition 11. *Assume unsuccessful negotiation attempts are unobservable. For any incentive compatible canonical contract profile without money burning in the original game there is a behavioral equivalent incentive compatible canonical contract profile in the game with endogenous negotiation attempts.*

Proof. Without money burning, all continuation payoffs of a canonical contract lie on the Pareto-frontier of PPE of the corresponding truncated game. This implies that in every state x and phase there is always at least one player who is weakly better off if new negotiations take place and another player who is weakly worse off. Let the contract profile in the game with endogenous negotiation be such that (at least) one player who weakly prefers new negotiation in state x attempts negotiations for state x and (at least) one player who does not prefer negotiation does not attempt negotiations; all other actions and transfers are the same as in the corresponding state and phase of the original game. It is straightforward that all relational contracts of that contract profile are incentive compatible and the contract profile is a canonical contract profile of the game with endogenous negotiation. \square

Sufficient condition that RNE negotiation payoffs can be implemented with canonical RNE

We now derive a sufficient condition on the stochastic game such that the negotiation payoffs of any RNE can always be implemented with a canonical RNE. This sufficient condition is satisfied by all examples in this paper. Games with monotone state transitions shall be stochastic games in which states cannot cycle in the following sense: if from a state x another state $x' \neq x$ can be reached with positive probability after some number of periods under some strategy profile then x can never be reached from state x' . Monotone state transitions imply that the game has at least one absorbing state, that can never change once reached. We say an action plan $(\alpha^k)_k$ is optimal given negotiation payoffs \mathbf{r} if the truncated game given \mathbf{r} has an optimal simple equilibrium with that action plan.¹⁷

Proposition 12. *For every RNE of a monotone stochastic game there exists a canonical RNE with the same negotiation payoffs.*

Proof. Consider some state $x \in X$ in a monotone stochastic game. The relational contracts chosen in state x do not affect the set of incentive compatible relational contracts in any state $x' \neq x$ that can be reached from x , because x cannot be reached from x' . Furthermore, the set of incentive compatible relational contracts in any state x'' that can reach x stays the same for all relational contracts in state x that yield the same expected negotiation payoffs. For a given RNE, we can thus replace the relational contracts of each player i in each state x by a relational contract that constitutes an optimal simple equilibria of the truncated game $\hat{\Gamma}(\mathbf{r}, x)$ and implements the original negotiation payoff \mathbf{r}^i without violating the incentive compatibility or optimality conditions of RNE. \square

Finding Optimal Simple Equilibria in Truncated Games

While there is no simple general recipe for finding canonical RNE, useful tools are results that facilitate computation of optimal simple equilibria in truncated games. For a truncated game with expected negotiation payoffs r , let $U(x|\alpha^e, r)$ denote the expected joint continuation payoffs (summing over payoffs for all players) in state x of a simple equilibrium with equilibrium phase policy α^e and no usage of money burning. These joint equilibrium payoffs can be easily computed by solving the following system of linear equations:

$$U(x|\alpha^e, r) = (1 - \delta)\Pi(\alpha^e, x) + \delta E[(1 - \rho)U(x'|\alpha^e, r) + \rho R(x')|\alpha^e, x], \quad (8)$$

where Π and R denote joint stage game payoffs and joint expected negotiation payoffs, respectively.

¹⁷It is straightforward that if the condition holds for the truncated game $\hat{\Gamma}(r, x)$ of some state $x \in X$ if and only if holds for the truncated games of all states $x \in X$.

The lowest punishment payoffs that can be imposed on player i given a punishment policy α^i are characterized by the solution $v_i(\cdot|\alpha^i, r)$ of the following Bellman equation:

$$v_i(x|\alpha^i, r) = \max_{\hat{a}_i \in A_i(x)} \{(1-\delta) (\pi_i(\hat{a}_i, \alpha_{-i}^i, x)) + \delta E[(1-\rho)v_i(x'|\alpha^i, r) + \rho r_i(x')|x, \hat{a}_i, \alpha_{-i}^i]\}. \quad (9)$$

It characterizes player i 's best-reply payoffs in the truncated game if we assume that other players' action plan is fixed to α_{-i}^i and no transfers are conducted. Adapting Proposition 5 from Kranz (2012), we find:

Proposition 13. *Consider a stochastic game with perfect monitoring and let $(\alpha^k)_k$ be an action plan in which in every state at least one player plays a pure strategy. A simple equilibrium with action plan $(\alpha^k)_k$ exists for the truncated game with negotiation payoffs r if and only if for every state $x \in X$ and every phase $k \in \{e, 1, \dots, n\}$*

$$(1-\delta)\Pi(\alpha^k, x) + \delta E[(1-\rho)U(x|\alpha^e, r) + \rho R(x')|\alpha^k, x] \geq \sum_{i=1}^n \max_{\hat{a}_i \in A_i(x)} \{(1-\delta)\pi_i(\hat{a}_i, \alpha_{-i}^k, x) + \delta E[(1-\rho)v_i(x'|\alpha^i, r) + \rho r_i(x')|\hat{a}_i, \alpha_{-i}^i, x]\}. \quad (\text{IC-SUM})$$

If (IC-SUM) is satisfied, the set of subgame perfect equilibrium payoffs that can be implemented in the truncated game with simple equilibria using action plan $(\alpha^k)_k$ is given by the simplex:

$$\{u \in \mathbb{R}^n \mid \sum u_i \leq U(x|\alpha^e, r) \text{ and } u_i \geq v_i(x|\alpha^i, r) \forall i\}$$

Similar results can be adapted for the general case of mixed strategies or for games with imperfect public monitoring of actions.

Appendix B: Proofs of Results in the Examples

Repeated Games

Proposition 3. *In a repeated game, negotiation payoffs are regular and their sum is equal to the highest joint public perfect equilibrium payoff $\bar{U}(\tilde{\delta})$ of the repeated game given an adjusted discount factor of*

$$\tilde{\delta} = (1-\rho)\delta.$$

Proof. The result is straightforward. For a given strategy profile σ , let $\tilde{\pi}(t, \sigma)$ denote the expected payoffs in period t in game with zero negotiation probability. The expected payoff in the truncated game $\Gamma(r, x)$ can be written as

$$u(\sigma) = \tilde{u}(\sigma) + r \sum_{t=1}^{\infty} \delta^t (1 - (1 - \rho)^t)$$

where

$$\tilde{u}(\sigma) = \sum_{t=0}^{\infty} (\delta(1 - \rho))^t \tilde{\pi}_t(t, \sigma)$$

is the payoff component corresponding to periods in which renegotiation has not yet taken place. The payoff function $\tilde{u}(\sigma)$ is identical to the discounted payoff of a repeated game with discount factor $\tilde{\delta} = (1 - \rho)\delta$. Since $u(\sigma)$ is simply $\tilde{u}(\sigma)$ plus a constant that does not depend on the strategy profile σ , the set of public perfect equilibria of the truncated game is the same as for the repeated game with discount factor $\tilde{\delta}$, independent of the negotiation payoffs r . This independence also implies that every player i will select a contract with regular negotiation payoffs. Regular negotiation payoffs are $\bar{U}(\tilde{\delta}) - \sum_{j \neq i} \bar{v}_j(\tilde{\delta})$ for player i and $\bar{v}_j(\tilde{\delta})$ for each player $j \neq i$ and have therefore a joint utility of $\bar{U}(\tilde{\delta})$. \square

Proposition 4. *In a repeated game with perfect monitoring expected negotiation payoffs of a canonical RNE with a pure action plan $(a^k)_k$ satisfy*

$$r_i = \pi_i^*(a^i) + \beta_i(\Pi(a^e) - \sum_{j=1}^n \pi_j^*(a^j)) \quad (10)$$

where $\pi^*(\alpha)$ denotes player i 's stage game best-reply payoff

$$\pi_i^*(\alpha) = \max_{a_i \in A_i} \pi_i(a_i, \alpha_{-i}).$$

Proof. Consider a canonical RNE with action plan $(a^k)_k$. We find from (8) that joint equilibrium phase payoffs and negotiation payoffs r satisfy in a repeated game

$$\begin{aligned} U^s &= (1 - \delta)\Pi(\mathbf{a}^e) + \delta((1 - \rho)U^s + \rho R) \\ &= \frac{1 - \delta}{1 - \tilde{\delta}}\Pi(\mathbf{a}^e) + \frac{\delta - \tilde{\delta}}{1 - \tilde{\delta}}R \end{aligned}$$

In a similar fashion it follows from (9) that player i 's punishment payoffs satisfy

$$\begin{aligned} v_i &= (1 - \delta)\pi_i^*(\mathbf{a}^i) + \delta((1 - \rho)v_i + \rho r_i) \\ &= \frac{1 - \delta}{1 - \tilde{\delta}}\pi_i^*(\mathbf{a}^i) + \frac{\delta - \tilde{\delta}}{1 - \tilde{\delta}}r_i \end{aligned}$$

Since negotiation payoffs are regular, they satisfy

$$\begin{aligned} r_i &= v_i + \beta_i \left(U - \sum_{j=1}^n v_j \right) \\ &= \frac{\delta - \tilde{\delta}}{1 - \tilde{\delta}} r_i + \frac{1 - \delta}{1 - \tilde{\delta}} \left(\pi_i^*(\alpha^i) + \beta_i \left(\Pi(\mathbf{a}^e) - \sum_{i=1}^N \pi_i^*(\alpha^i) \right) \right). \end{aligned}$$

Solving for r_i yields

$$r_i = \pi_i^*(\alpha^i) + \beta_i \left(\Pi(\mathbf{a}^e) - \sum_{i=1}^N \pi_i^*(\alpha^i) \right).$$

□

Principal-Agent Game with Endogenous Inside Option

Proposition 5. *If a lower inside option is socially more efficient in the repeated game, i.e. $\tilde{U}(x_L) > \tilde{U}(x_H)$, then in every Pareto-optimal subgame perfect equilibrium the principal destroys her inside option.*

Proof. We first note that minimal punishment payoffs in state x satisfy $\bar{v}_1(x) = x$ for the principal and $\bar{v}_2(x) = 0$ for the agent. Given that the SPE payoff set is the simplex described in Theorem 1, it suffices to show that $\tilde{U}(x_L) > \tilde{U}(x_H)$ implies that a higher joint payoff can be implemented in state x_H if the inside option is immediately destroyed. For a proof by contradiction assume $\tilde{U}(x_L) > \tilde{U}(x_H)$ holds but it is Pareto-optimal not to destroy the inside option. Highest joint payoffs SPE are then given by $\bar{U}(x_H) = \tilde{U}(x_H)$. Since $\bar{U}(x_L) = \tilde{U}(x_L)$ this implies $\bar{U}(x_H) < \bar{U}(x_L)$. Consider the joint incentive constraint from Proposition 13 for implementing effort e_H in state x_H with zero negotiation probability. If the inside option is destroyed, it is

$$\delta \bar{U}(x_L) \geq (1 - \delta)k(e_H) + \delta x_H$$

and otherwise

$$\delta \bar{U}(x_H) \geq (1 - \delta)k(e_H) + \delta x_H.$$

Since $\bar{U}(x_H) < \bar{U}(x_L)$, higher effort levels and thus higher joint payoffs can be implemented in state x_H if the inside option is simultaneously destroyed, contradicting the assumption that it is Pareto-optimal not to destroy the inside option. □

Preliminary observations for the proofs of Propositions 6 and 7 We will omit some straightforward but tedious case distinctions, and assume henceforth that it is a best-reply for the principal not to destroy the inside option when

punished, the agent chooses zero effort in the principal's punishment phase and the principal does not destroy her inside option in the agent's punishment phase.

Once state x_L is reached, players face a repeated game with negotiation payoffs $\tilde{r}(x_L)$ and maximal joint payoffs $\tilde{U}(x_L)$. We derive a convenient form of the incentives constraints for the equilibrium phase in state x_H for simple relational contracts of the truncated game for the two cases that the inside option is destroyed or not.

First consider simple equilibria, in which in the equilibrium phase in state x_H the principal destroys her inside option and effort e_H is chosen; in state x_L a joint repeated game payoff of $\tilde{U}(x_L)$ is implemented. For all levels of e_H , the joint incentive constraint (SUM-IC) can be reformulated as

$$\tilde{\delta} \left(\tilde{U}(x_L) - x_H \right) - (1 - \tilde{\delta})k(e_H) \geq \frac{\omega}{1 - \omega} (r_1(x_H) - \tilde{r}_1(x_L)). \quad (11)$$

with

$$\omega \equiv \frac{\delta - \tilde{\delta}}{1 - \tilde{\delta}} \in [0, 1).$$

The parameter ω can be interpreted as the relative share of adjusted discounting $(1 - \tilde{\delta})$, that is due to the negotiation probability ρ , while $1 - \omega = \frac{1 - \delta}{1 - \tilde{\delta}}$ can be interpreted as the relative share explained by the discount factor δ . When $\tilde{\delta}$ is fixed, a larger ω is equivalent to a larger negotiation probability ρ .

The joint incentive constraint for a simple contract in which the outside option is not destroyed and maximum joint payoffs $\bar{U}(x_H)$ are implemented can be reformulated as

$$\tilde{\delta} \left(\bar{U}(x_H) - x_H \right) - (1 - \tilde{\delta})k(e_H) \geq 0. \quad (12)$$

We now prove

Proposition 6. *Holding the adjusted discount factor $\tilde{\delta}$ fixed, the principal destroys her inside option in a RNE if and only if $\bar{U}(x_H) \leq \tilde{U}(x_L)$ and the negotiation probability is below a critical value $\bar{\rho}(\beta_1)$ that increases in the principal's bargaining weight β_1 .*

Proof. It is straightforward to show that if $\bar{U}(x_H) > \tilde{U}(x_L)$ then the inside option will never be destroyed in a RNE. We thus assume that $\bar{U}(x_H) \leq \tilde{U}(x_L)$. Recall that negotiation payoffs in the repeated game with fixed inside option are given by

$$\tilde{r}_1(x) = (1 - \beta_1)x + \beta_1\tilde{U}(x).$$

Whenever the principal's bargaining weight β_1 exceeds the critical value

$$\beta_1^* \equiv \frac{x_H - x_L}{\left(\tilde{U}(x_L) - x_L \right) - \left(\tilde{U}(x_H) - x_H \right)} \in (0, 1)$$

we have $\tilde{r}_1(x_H) \leq \tilde{r}_1(x_L)$, and it is in both the principal's and agent's interest to select a contract in which the inside option is destroyed.

Consider now the case $\tilde{r}_1(x_H) \geq \tilde{r}_1(x_L)$. It always holds true that $r_1(x_H) \geq \tilde{r}_1(x_H)$ and the incentive constraint for destroying the inside option is easiest satisfied for $r_1(x_H) = \tilde{r}_1(x_H)$. For this value of negotiation payoffs the incentive constraint for destroying the inside option (11) when choosing effort e_H in state x_H becomes

$$\tilde{\delta} (\bar{U}(x_L) - x_H) - (1 - \tilde{\delta})k(e_H) \geq \frac{\omega}{1 - \omega} (x_H - x_L - \beta_1 ((\bar{U}(x_L) - x_L) - (\bar{U}(x_H) - x_H))).$$

If $\beta_1 < \beta_1^*$ the term in the brackets on the right hand side is positive and for any fixed level of e_H for which the left hand side is positive, there exist a maximum value of $\bar{\omega}(\beta_1) \in (0, 1)$ for which this condition can be satisfied. The result follows, because a contract in which the inside option is destroyed is feasible only if the condition can be satisfied for $e_H = 0$ and desirable only if the condition can be satisfied for sufficiently high values of e_H . \square

Proposition 7. *Fix a negotiation probability $\rho \in (0, 1)$ and consider the limit $\delta \rightarrow 1$. The principal destroys her inside option if and only if her negotiation payoff in the repeated game with fixed inside option is larger under the low inside option, i.e. $\tilde{r}_1(x_H) \leq \tilde{r}_1(x_L)$.*

Proof. For the case $\tilde{U}(x_L) \leq x_H$ it is straightforward that $\tilde{r}_1(x_H) \geq \tilde{r}_1(x_L)$ and that the principal will not destroy the inside option in any RNE. Assume therefore $\tilde{U}(x_L) \geq x_H$, which implies $\tilde{U}(x_L) > \tilde{U}(x_H)$. We see from (11) that the destruction of the inside option can be easiest implemented if the agent chooses zero quality in the first period, i.e. $q_H = 0$. Intuitively, this relaxes the agent's incentives to deviate as much as possible and allows the largest incentive compatible reward payment from agent to principal for destroying the inside option. For $q_H = 0$, the destruction of the inside option can be implemented if and only if the principal's negotiation payoff satisfies

$$r_1(x_H) \leq \tilde{r}_1(x_L) + \varepsilon(\omega)$$

where $\varepsilon(\omega)$ is a strictly positive function with $\lim_{\omega \rightarrow 1} \varepsilon(\omega) = 0$. The limit of $\omega \rightarrow 1$ corresponds to the case that for a fixed negotiation probability $\rho > 0$ the discount factor δ converges towards 1. Intuitively, in this limit of $\delta \rightarrow 1$, continuation payoffs will always be approximately equal to negotiation payoffs and the principal will pick the state with the higher negotiation payoffs. Since the principal has the option not to destroy the inside option, she can guarantee herself the same negotiation payoff as if the state were fixed to x_H , i.e.

$$r_1(x_H) \geq \tilde{r}_1(x_H).$$

The two inequalities imply that the principal will not destroy her inside option if $\tilde{r}_1(x_L) < \tilde{r}_1(x_H)$.

Now consider the case $\tilde{r}_1(x_L) > \tilde{r}_1(x_H)$. Consider the limit $\delta \rightarrow 1$ and first consider the case that there would exist a RNE with immediate destruction of the inside option that has regular negotiation payoffs. These regular negotiation payoffs then satisfy $r_1(x_H) > \tilde{r}_1(x_L)$ because the joint payoff in state x_H is approximately $\tilde{U}(x_L)$ but the principal gets a strictly higher share in state x_H since an initial offer by the agent must grant the principal at least x_H . Such regular payoffs can, of course, not be implemented for sufficiently large ω , since they would violate the limit condition $r_1(x_H) \leq \tilde{r}_1(x_L)$, i.e. the principal could not be incentivized to destroy her inside option. Yet, if $\tilde{r}_1(x_L) > \tilde{r}_1(x_H)$, the principal can and wants to offer (irregular) relational contracts in which the inside option is destroyed that grant her expected negotiation payoffs $r_1(x_H)$ satisfying $\tilde{r}_1(x_H) < r_1(x_H) \approx \tilde{r}_1(x_L)$. In a related fashion, one can also show that for the knife-edge case $\tilde{r}_1(x_L) = \tilde{r}_1(x_H)$, both destruction and non-destruction can be implemented in a RNE. \square

Inside Options vs Outside Options

The proof follows straightforward from the arguments given in the text.

A simple Arms Race

Proposition 9. *In the unique Markov perfect equilibrium outcome, as well as, in all Pareto-optimal subgame perfect equilibrium outcomes no weapons are bought or used; this also holds true for the corresponding RNE in the extreme cases of renegotiation in every period ($\rho = 1$) or no renegotiation ($\rho = 0$). In contrast, for intermediate renegotiation probabilities, it can be the case that there is a unique RNE outcome in which one or both players build weapons.*

Proof. In a Markov perfect equilibrium players will not conduct transfers nor use weapons, since both reduce own payoffs without affect the state, i.e. they cannot have an effect on future play. It follows immediately that no player will spend the cost to buy a weapon in a MPE. Hence, the MPE is unique and implements the highest feasible joint payoff of zero. If we have renegotiation in every period ($\rho = 1$), continuation play only depends on the current state and the unique RNE is equal to the unique MPE. It follows from the linear Pareto-frontier of subgame perfect equilibrium payoffs, that also every Pareto-optimal optimal SPE equilibrium, and the corresponding RNE given no renegotiation ($\rho = 0$) must implement a joint payoff of zero.

Consider now the case of intermediate renegotiation probabilities. To illustrate that there can exist RNE in which one player builds weapons, consider the simple variant that only player 1 has the opportunity to build a weapon and that this opportunity only occurs in period 1. From period 2 onwards the state is fixed and players essentially play a repeated game.

Consider this repeated game for the case that player 1 has bought a weapon in period 1. Punishment of player 2 is punished by player 1 using the weapon can be implemented in this repeated game, whenever the following joint incentive constraint (see Proposition 13) is satisfied:

$$(1 - \tilde{\delta})\varepsilon \leq \tilde{\delta}d$$

Cost of weapon usage must be sufficiently small compared to the induced damage, and the adjusted discount factor cannot be too high. If that condition is satisfied, then if renegotiation take place in period 2 or later, player 1 select a relational contract that in which player 2 pays a transfer of d on the equilibrium phase in order to avoid punishment with the weapon. Player 2 obviously selects a contract in which he does not transfer anything. Expected negotiation payoffs of player 1 are then given by $\beta_1 d$.

By rearranging the joint incentive constraint (IC-SUM), we find that there exists no relational contract in period 1 that prevents player 1 from buying a weapon whenever

$$(1 - \delta)b < \omega\beta_1 d.$$

If the negotiation probability is zero ($\omega = 0$) then this condition is always satisfied: one simply chooses continuation equilibria in which transfers are never conducted, irrespective of the potential damage that a player with weapons can impose. As long as it is credible to use weapons, a higher negotiation probability makes it harder to prevent acquisition of weapons since players can earlier monetize their potential to harm the other player. Once the negotiation probability is too large, it is no more credible to use weapons, however. \square