DISCUSSION PAPER SERIES

DP12540

RECONCILIATING RELATIONAL CONTRACTING AND HOLD-UP: A MODEL OF REPEATED NEGOTIATIONS

Susanne Goldlücke and Sebastian Kranz

INDUSTRIAL ORGANIZATION

ISSN 0265-8003

RECONCILIATING RELATIONAL CONTRACTING AND HOLD-UP: A MODEL OF REPEATED NEGOTIATIONS

Susanne Goldlücke and Sebastian Kranz

Discussion Paper DP12540
Published 24 December 2017
Submitted 24 December 2017

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

RECONCILIATING RELATIONAL CONTRACTING AND HOLD-UP: A MODEL OF REPEATED NEGOTIATIONS

Abstract

We propose a unified framework to study relational contracting and hold-up problems in infinite horizon stochastic games with monetary transfers. Starting from the observation that 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 negotiated in a relationship. New negotiations take place with positive probability each period and treat previous informal agreements as bygones. The concept nests relational contracting and hold-up models as opposite corner cases. Allowing for intermediate cases sheds light on many plausible trade-offs that do not arise in these corner cases.

JEL Classification: C73, C78, D23, L14

Keywords: relational contracting, hold-up, negotiations, Stochastic Games

Susanne Goldlücke - susanne.goldluecke@uni-konstanz.de
University of Konstanz and CEPR

Sebastian Kranz - sebastian.kranz@uni-ulm.de
University of Ulm

Acknowledgements
We would like to thank Daniel Barron, Mehmet Ekmekci, Eduardo Faingold, Matthias Fahn, Paul Heidhues, Johannes Hörner, David Miller, Larry Samuelson, Klaus Schmidt, Patrick Schmitz, Philipp Strack, Juuso Välimäki, Nicolas Vieille, Joel Watson and many seminar participants for very helpful discussions on this and earlier versions of this paper.

Reconciliating Relational Contracting and Hold-up: A Model of Repeated Negotiations*

Susanne Goldlücke† and Sebastian Kranz‡

December 22, 2017

Abstract

We propose a unified framework to study relational contracting and hold-up problems in infinite horizon stochastic games with monetary transfers. Starting from the observation that 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 negotiated in a relationship. New negotiations take place with positive probability each period and treat previous informal agreements as bygones. The concept nests relational contracting and hold-up models as opposite corner cases. Allowing for intermediate cases sheds light on many plausible trade-offs that do not arise in these corner cases.

Keywords: relational contracting, hold-up, negotiations, stochastic games
JEL codes: C73, C78, D23, L14

1 Introduction

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

*We would like to thank Daniel Barron, Mehmet Ekmekci, Eduardo Faingold, Matthias Fahn, Paul Heidhues, Johannes Hörner, David Miller, Larry Samuelson, Klaus Schmidt, Patrick Schmitz, Philipp Strack, Juuso Välimäki, Nicolas Vieille, Joel Watson and many seminar participants for very helpful discussions on this and earlier versions of this paper.

†Department of Economics, University of Konstanz. Email: susanne.goldluecke@uni-konstanz.de. Phone: +497531882989

‡Department of Mathematics and Economics, Ulm University. Email: sebastian.kranz@uni-ulm.de. Phone: +497315023691

possibly affect the way how that surplus is distributed. Limitations to formal contracting in such relationships have inspired two large branches of economic literature. First, there is the literature on the hold-up problem, which occurs if long-term investments cannot be protected by complete contracts.[1] Second, there is the literature on relational contracts, which use repeated interaction and credible punishments to enforce 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 perfect public equilibria (PPE) of infinitely repeated games with payoffs on the Pareto frontier, allowing for monetary transfers between the players. In order to model relationships with long-term investments and corresponding hold-up problems, we study dynamic stochastic games, in which the stage game can change over time in response to players' actions. As we will illustrate with a simple example, relational contracts can easily overcome many hold-up problems by making future trade and bargaining outcomes dependent on the conducted investments. This means that incomplete formal contracting is not sufficient for the existence of hold-up problems. It is also crucial that relational contracting is incomplete such that to a certain extent bygones are treated as bygones. Being able to account for this driving force of hold-up problems in models of relational contracting is the key motivation for introducing our concept of *repeated negotiation equilibrium* (RNE).

An extreme form of incomplete relational contracting would be that in each period, continuation play is completely determined by new negotiations. This idea is opposite to an essential feature of relational contracting, namely the history-dependence of continuation play. Our model of repeated negotiations 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 aspect of the history.

[1] The hold-up problem has received a lot of attention since Grout's (1984) classical article, which shows how firms under-invest in capital because labor unions appropriate a share of the generated surplus in subsequent wage negotiations. Investment inefficiencies and the interaction with negotiation outcomes lie at the heart of hold-up problems, which affect the way that production is organized (Klein et. al. (1978), Williamson (1985), Hart and Moore (1988)).

[2] Self-enforcing contracts between firms and employees or between firms and their suppliers are for example studied in Bull (1987), MacLeod and Malcomson (1989), Levin ((2002), (2003)), Baker, Gibbons and Murphy (2002), and Board (2011). See Malcomson (2012) for a survey.

That new negotiations are triggered by sun-spot events has certain intuitive appeal and is a simple way to introduce a continuous measure of the importance of history-independent bargaining power. A larger negotiation probability means that expected payoffs have to return sooner to history-independent bargaining payoffs. In the corner case of a negotiation probability of zero, players can commit to any credible path of play such that an RNE corresponds to a Pareto-optimal PPE. If the negotiation probability is one and the game has a unique Markov perfect equilibrium (MPE), then the RNE corresponds to that MPE.

A requirement for an RNE is that a player selects his most preferred relational contract among those which are incentive compatible given future negotiations and given the belief that the player would choose that relational contract again in future negotiations in the same state. This assumption captures a realistic aspect of negotiations in repeated relationships and ensures existence of RNE. Our existence theorem also shows that there always exist RNE with a simple, tractable form: All relational contracts have a stationary structure on the equilibrium path and negotiations affect the path of play only by changing the transfers that take place following negotiations.

In the special case of an infinitely repeated game, the predictions of RNE are not qualitatively different from Pareto-optimal PPE. In repeated games, actions have no payoff-relevant long run effects and negotiation outcomes are therefore not affected by past decisions. The exogenous negotiation probability then has the same effect as including a probability with which the relationship ends after each period. Negotiation can be interpreted as a restart of the relationship, and a positive negotiation probability simply adjusts the discount factor downwards.

This equivalence between lower discount factors and higher negotiation probabilities no longer holds if actions have long-term effects. Applying the concept of RNE in such settings allows to meaningfully study the interaction between relational contracting and hold-up. To illustrate how repeated negotiations affect implausible predictions of Pareto-optimal PPE, we analyze a principal-agent relationship in which the principal can make herself permanently more vulnerable at zero cost. In a Pareto-optimal PPE, the principal would always make use of this possibility. There is no drawback for the principal since Pareto-optimal PPE allow to perfectly coordinate away from any undesired abuse of the created vulnerabilities.[3] In contrast, in a hold-up context parties

[3]This is related to the effect that institutions that are inferior in a stage game may be preferred once the game is repeated (e.g. Baker, Gibbons and Murphy (1994), Schmidt and Schnitzer (1995), Pearce and Stacchetti (1998), Halonen (2002), Iossa and Spagnolo (2011)). While institutions are often

would never attempt to weaken their bargaining position. We refer to this contradiction as the “Vulnerability Paradox”. With an intermediate negotiation probability, the principal solves a natural trade-off between these two forces. For a similar reason, repeated negotiations cause parties to pay for costly compliance systems or to use a gradual process to guarantee mutual vulnerability and similar bargaining positions.

Repeated negotiations may also create incentives to make others more vulnerable: In another example we illustrate how the concern about future negotiations can induce costly arms races even when raising arms against other players involves costs but no direct gains. Finally, two further examples shed light on the negotiation payoffs in an RNE. In a repeated game, the negotiation payoffs of a RNE split the surplus as in a generalized Nash bargaining solution with disagreement points equal to the lowest possible payoff that players can get in a relational contract. We explicitly introduce inside and outside options in a repeated principal-agent model and find that the outside option principle is consistent with our model of repeated negotiations. In general, negotiation payoffs in a stochastic game need not have this Nash bargaining form, as we illustrate with a blackmailing game. In this example, we show how repeated negotiations render blackmailing threats incredible and analyze when brinkmanship can make extortion possible.

We are only aware of a few papers that have studied the interaction of investments, hold-up, and relational contracting. Garvey (1995), Baker et al. (2002), Halonen (2002), and Blonski 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 (2015) 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. Our results contribute to 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.[4]

The general idea of RNE is most closely related to Miller and Watson’s (2013, MW)

---

exogenous and affect all players in the same way, the players will also always have an incentive to select an inferior institution, even if it puts some players at a disadvantage in a one-shot interaction.

[4]Outside the relational contracting literature, a few papers have studied how dynamic interaction affects the hold-up problem. For example, Aghion, Dewatripont and Rey (1994) show how the design of ex post renegotiation in an infinitely repeated bargaining game can solve a static hold-up problem. Che and Sákovics (2004) and Pitchford and Snyder (2004) study hold-up problems in stochastic games with sequential investment decisions. They assume that once investment stops, the resulting surplus is split via an enforceable contract.

concept of contract equilibria for repeated games. Both concepts interpret relational contracting as a process of repeated negotiations over continuation play. MW assume that new negotiations take place in every period and consider a negotiation procedure with an explicit disagreement point, which can itself depend on the history. The main factor by which negotiations introduce inflexibility to relational contracting in MW is that in periods of disagreement, players cannot conduct transfers to each other. In contrast, the random dictator procedure in our framework does not specify disagreement points. With the negotiation probability we add a degree of freedom relative to MW, which however is a useful parameter to study comparative statics of relational contracting once stochastic games with long-term decisions are considered.

The idea that relational contracts can be renegotiated during the relationship has also been explored in the literature on renegotiation-proofness in repeated games, e.g. Bernheim and Ray (1989), Farrell and Maskin (1989), Asheim (1991), Abreu, Pearce and Stacchetti (1993), and Safronov and Strulovici (2016). 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 equilibrium 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.

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 the general formulation and characterization of repeated negotiation equilibria. Section 4 illustrates the concept for several relational contracting examples with long run decisions. Proofs are relegated to an appendix.

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 $\delta \in (0,1)$.

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

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

The hold-up problem arises if contracts are incomplete and the parties bargain over trade decision and distribution of the surplus only once investment decisions are made. In the hold-up literature it is commonly assumed that surplus from trade is split according to the (symmetric) Nash bargaining solution, which in our example corresponds to an equal split of $S(x)$. Hence no player receives the full return to their investment and incentives to invest are distorted. It is in general not possible to implement both first best investments and ex-post efficient trading decisions with simple contracts in this setting. Note that the model allows for cooperative investments, i.e. the seller's investments can influence the buyer's valuation and vice versa.[5]

The following result states the straightforward observation that if we remove the assumption that the surplus is split according to the Nash bargaining solution, the hold up problem can be overcome.

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

To show this, we assume that first best investments a^* are strictly positive for at least one player, since 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

[5] Che and Hausch (1999) show that with unobservable cooperative investments, the hold-up problem cannot be resolved by any contract. If 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 with simple enforceable contracts that act as a threatpoint in the renegotiations (e.g. Nöldeke and Schmidt (1995) and Edlin and Reichelstein (1996)).

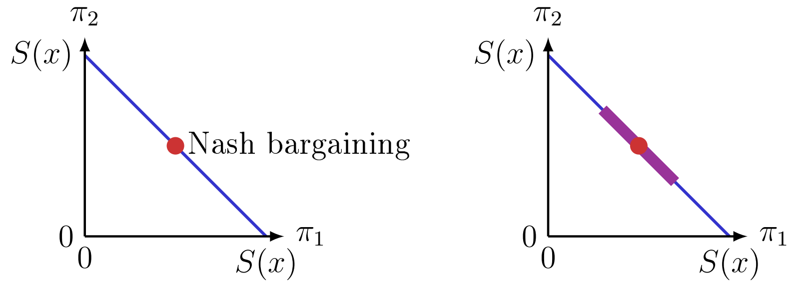

Figure 1: Set of continuation payoffs in period 2. The thick line segment in the right figure illustrates the range of expected continuation payoffs that can be implemented for a negotiation probability of $\rho = 0.6$.

0 in all states. If no player has unilaterally deviated, we pick continuation equilibria that split the surpluses $S(x)$ such that on average 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.

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 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.[6] As Evans (2008) shows, the result that the hold-up problem can be overcome with simple formal contracts when the bargaining game has multiple continuation equilibria holds very generally, with potentially infinite bargaining games and unobservable, possibly cooperative, investment.

Imposing the Nash bargaining solution corresponds to the idea that previous non-enforceable agreements on how to split trade surplus are ignored once investment costs are sunk, which is a cornerstone of the hold-up literature. In contrast, relational contracts are built around the idea that past non-enforceable agreements always remain valid. 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

[6]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. The famous exception is the unique implementation of the Nash bargaining outcome in the alternative offer bargaining game by Rubinstein (1982), which however is not robust with respect to plausible modifications of the bargaining game (e.g. Avery and Zemsky (1994)).

intermediate cases.[7]

In the example, a natural formulation of intermediate cases would be to require 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 1 (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 some probability to select the new relational contract and then 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 on negotiation taking 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.[8]

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 1 (right) thus shows the range of implementable expected payoffs for $\rho = 0.6$.

[7] See Ellingsen and Robles (2002) and Tröger (2002) for evolutionary arguments on appropriate equilibrium selection. Ellingsen and Johannesson (2004) investigate hold-up problems experimentally, in light of the observation that the Nash demand game allows more SPE than an ultimatum bargaining game. Their results support the view that intermediate cases are plausible.

[8] Since 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.

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

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

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

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

In line with Proposition 1, we find that absent repeated negotiation ($\rho = 0$) players always implement first best investments. 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. The reason is that a lower discount factor also makes high investments levels less desirable from a social perspective.[9]

In the limit case of no discounting $\delta \to 1$, 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

[9]For 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 ambiguos 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.

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 increases in the parameter γ that determines the gross social 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.

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 alone does not restrict the the ability to implement first best investments in the example since the state x is sufficiently informative about the investment decisions. Repeated negotiation equilibria for the case $\rho = 1$ are equivalent to specific MPE that imply a strong notion of bygones.

3 Repeated Negotiation 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.

In the transfer stage, every player simultaneously chooses a non-negative vector of transfers to all other players.[10] 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 $A(x) = A_1(x) \times ... \times A_n(x)$.

After actions have been conducted, a signal y from a finite signal space Y and a new

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

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 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 payoffs, i.e. the expected sum of 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 (SPE).

3.2 Repeated Negotiation Equilibria

For our concept of a repeated negotiation equilibrium, we assume in addition that at the beginning of each period t, players commonly observe a public signal R_t that determines whether there are new negotiations. It can take values in $\{0, 1, .., n\}$, where

$R_t = 0$ occurs with probability $1-\rho$, and $R_t = i \in \{1, ..., n\}$ occurs with probability $\rho\beta_i$ for some $\rho, \beta_1, ..., \beta_n \in [0,1], \sum_{i=1}^n \beta_i = 1$. The parameter ρ is called the negotiation probability, since $R_t \neq 0$ indicates that new negotiations take place. It is also assumed that R_0 takes on values in $\{1, .., n\}$ only, with $\text{Prob}[R_0 = i] = \beta_i$, which means that in the first period, negotiations always take place.

A relational contract shall be an incomplete strategy profile that describes play just until new negotiations take place. Denoting by $H^0(x)$ the set of all histories beginning with state x in the game without negotiations (i.e., histories with all $R_t = 0$), a relational contract maps histories in $H^0(x)$ to actions or transfers. Repeated negotiations shall follow a simple random dictator procedure. At the beginning of each period t, with a negotiation probability ρ, the sunspot signal R_t takes on values in $\{1, ..., n\}$, indicating which player can select a relational contract. The parameter β_i is the probability that player i is chosen 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. It cannot condition on any event that occurred before the negotiations. A helpful image is that players forget the history of play when negotiations take place and only remember the current state x.

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

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

is called a relational contract profile. Every relational contract profile constitutes a strategy profile of the stochastic game with transfers and negotiation signals. Formally, this strategy profile is defined by $\sigma(h) = \sigma_{(i,x)}(\tilde{h})$, where $h = (..., x, i, \tilde{h})$ and $\tilde{h} \in H^0(x)$. A relational contract profile has the property that continuation strategies are equal for all histories that end in the same state and with the same realization of the negotiation signal. We denote by $\sigma_{-(i,x)}$ a relational contract profile that excludes the relational contract selected by player i in state x.

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

$$r(x|\sigma) = \sum_{i=1}^{n} \beta_i u(\sigma|x, i) \tag{2}$$

denote the expected continuation payoffs in state x, if it were known that renegotiation takes place but not which player can select the new relational contract.

For the definition of repeated negotiation equilibrium we introduce a class of stochastic games we call *truncated games.* We refer to a function r that maps every state x into a payoff vector in the compact and convex set of feasible payoff vectors in state x as negotiation payoffs. A *truncated game* $\Gamma(r, x_s)$ is parameterized by arbitrary negotiation payoffs r (for each state) and an initial state $x_s \in X$. Compared to the original game, state transitions are modified such that if the original game transits to state x, it now only does so with probability $1-\rho$. With probability ρ, an absorbing state x_r is reached instead in which players receive $r(x)$ in every future period. The truncated game starts in a non-absorbing state x_s, such that there is at least one period of play before an absorbing state is reached. As long as no absorbing state is reached, payoffs and action spaces of the truncated game are the same as in the original game. Any truncated game is again a stochastic game with transfers, and relational contracts can be interpreted as strategy profiles of truncated games. The definitions directly imply that a relational contract profile σ constitutes a PPE of the original game if and only if for every player i and every state x_s, the relational contract $\sigma_{(i,x_s)}$ constitutes a PPE of the truncated game $\Gamma(r(.|\sigma), x_s)$. For a given contract profile $\sigma = (\sigma_{(i,x)})_{i\in\{1,\ldots,n\},x\in X}$ and arbitrary negotiation payoffs r, let $u(\sigma_{(i,x_s)}|r)$ denote the players' expected payoff from $\sigma_{(i,x_s)}$ in the truncated game $\Gamma(r, x_s)$ and define $g(x_s|r,\sigma) = \sum_{i=1}^{n} \beta_i u(\sigma_{(i,x_s)}|r)$.

Lemma 1. *The fixed point equation $r = g(.|r,\sigma)$ has the unique solution $r = r(.|\sigma)$.*

Consider a relational contract $\sigma_{(i,x)}$ chosen by player i in state x and take as given a profile of other contracts $\sigma_{-(i,x)}$. We say $\sigma_{(i,x)}$ is *incentive compatible* given $\sigma_{-(i,x)}$ if $\sigma_{(i,x)}$ is a PPE in the truncated game $\Gamma(r(.|\sigma), x)$, i.e. if no player has an incentive to deviate from $\sigma_{(i,x)}$ if player i always selects it in state x.

Intuitively, in a repeated negotiation equilibrium each random dictator will choose in negotiations an incentive compatible relational contract that maximizes her expected payoffs taking into account future negotiations. However, for our existence theorem, we need a slightly weaker concept in which a random dictator can ignore better contracts if they are not *stable* in the following sense. A relational contact $\sigma_{(i,x)}$ is *stable incentive compatible* given $\sigma_{-(i,x)}$ if for all sufficiently small modifications to $\sigma_{-(i,x)}$ we can find a relational contract that is incentive compatible given the modified profile and close to $\sigma_{(i,x)}$. Formally, a relational contract $\sigma_{(i,x)}$ is stable incentive compatible given $\sigma_{-(i,x)}$ if for all sequences $\sigma^m_{-(i,x)}$ that converges pointwise to $\sigma_{-(i,x)}$, there exists a subsequence $\sigma^{m_k}_{-(i,x)}$ such that for all $k \in \mathbb{N}$ there exists a PPE $\sigma^k_{(i,x)}$ in the truncated game $\Gamma(r(\sigma^k_{(i,x)}, \sigma^{m_k}_{-(i,x)}), x)$ such that $\sigma^k_{(i,x)}$ converges pointwise to $\sigma_{(i,x)}$.

Definition 1. A relational contract profile $\sigma = (\sigma_{(i,x)})_{i \in \{1,\ldots,n\}, x \in X}$ is called a *repeated negotiation equilibrium (RNE)* with expected payoffs $r(x^0|\sigma)$ 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 stable incentive compatible given $\sigma_{-(i,x)}$ and strictly preferred by player i, i.e., that satisfies $u_i(\tilde{\sigma}_{(i,x)}, \sigma_{-(i,x)}|x, i) > u_i(\sigma|x, i)$.

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 the idea that players neglect 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.[11]

As already noted, the qualifier "stable" for the incentive compatibility of contracts that the RNE is compared to is inserted for technical reasons only in order to ensure existence. In all our applications, we use the slightly stronger definition of an RNE without this qualifier. Further scope for simplification arises from the fact that every PPE payoff of a truncated game can be implemented with an optimal simple equilibrium (Goldlücke and Kranz (2017)).[12] Indeed, the condition in the definition of an RNE only needs to be checked for simple equilibria.

Lemma 2. *A relational contract profile σ is a RNE if for every state x and every player i, the relational contract $\sigma_{(i,x)}$ is incentive compatible given $\sigma_{-(i,x)}$ and there exists no simple relational contract $\tilde{\sigma}_{(i,x)}$ that is incentive compatible given $\sigma_{-(i,x)}$ with $u_i(\tilde{\sigma}_{(i,x)}, \sigma_{-(i,x)}|x, i) > u_i(\sigma|x, i)$.*

In a truncated game with negotiation payoffs r, let $\bar{U}(x_s|r)$ denote the maximum of the joint PPE continuation payoffs at the beginning of a period in state x_s and $\bar{v}_i(x_s|r)$

[11] 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.

[12] See Appendix B for the definition of a simple strategy profile.

the corresponding minimum of player i's PPE continuation payoffs. If it holds that expected negotiation payoffs are

$$r_i(x) = \bar{v}_i(x|r) + \beta_i(\bar{U}(x|r) - \sum_{j=1}^{n} \bar{v}_j(x|r)), \tag{3}$$

then we say that they are *regular*. With regular negotiation payoffs, the players split the highest joint continuation payoff of the truncated game 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 in the truncated game. If an RNE has regular payoffs, the value of ρ measures how close continuation payoffs have to be to what payoffs would be if they were negotiated as in the first period. It is not always the case that RNE have regular negotiation payoffs, as the blackmailing game in Subsection 4.5 illustrates.

3.3 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 $r(.|\sigma)$; if σ is also a RNE, we call it a canonical RNE. Since negotiations affect the path of play only by modifying the subsequent upfront payments, a canonical RNE has a particularly simple and tractable structure.

Theorem 1. *If the action space is finite and mixed actions are allowed then a canonical RNE exists.*

Since canonical RNE always exist, it may make sense to restrict attention to canonical RNE. One can show that for every RNE σ there exists an incentive compatible canonical contract profile $\tilde{\sigma}$ that has the same negotiation payoffs. This suggests that there is little lost by restricting attention to canonical contract profiles. 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).

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. 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.

Proposition 1. *For every RNE of a stochastic game with monotone state transitions there exists a canonical RNE with the same negotiation payoffs.*

3.4 Repeated games

A repeated game with transfers corresponds to the special case that there is just a single state. The set of PPE payoffs of a repeated game for given discount factor is given by $\{u \in \mathbb{R}^n | \sum u_i \leq \bar{U}(\delta) \text{ and } u_i \geq \bar{v}_i(\delta) \text{ for all } i\}$, where $\bar{U}(\delta)$ is the highest joint PPE payoff and $\bar{v}_i(\delta)$ is the lowest PPE payoff for player i.[13] The concept of a repeated negotiation equilibrium does not significantly change the analysis of repeated games. Basically, a positive negotiation probability just reduces the effective discount factor and the bargaining weights pin down which payoff vector on the Pareto frontier of PPE payoffs is selected.

Proposition 2. *In a repeated game, negotiation payoffs are regular and their sum is equal to the highest joint PPE payoff of a repeated game given an adjusted discount factor of $\tilde{\delta} = (1-\rho)\delta$. Expected negotiation payoffs satisfy*

$$r_i = \bar{v}_i(\tilde{\delta}) + \beta_i(\bar{U}(\tilde{\delta}) - \sum_{j=1}^{n} \bar{v}_j(\tilde{\delta})). \tag{4}$$

This result reflects the fact that in repeated games, a probability that the relationship ends is equivalent to a lower discount factor. New negotiations essentially constitute a termination and restart of the relationship.

In applications of repeated games, critical discount factors are often used to compare institutions with respect to the ability to sustain first-best outcomes in a relational contract. We have just shown that in repeated games critical negotiation probabilities are basically equivalent to critical discount factors. This result is reassuring since it

[13] See Goldlücke and Kranz (2012) for details.

means that the usual analysis of institutions and relational contracting is robust to the introduction of repeated negotiations.

When it comes to stochastic games, first-best strategies may depend on the discount factor, so that it has anyway little appeal to study the minimal discount factors under which first-best strategies can be implemented. In contrast, the first-best is not affected by the negotiation probability, which makes critical negotiation probabilities more suitable measures to study comparative statics of relational contracting in stochastic games.

4 Examples

This section illustrates the effects of repeated negotiation in relational contracting with simple examples, which all satisfy the sufficient condition in Proposition 1.

4.1 The Vulnerability Paradox in a Principal-Agent Relationship

Our first example is a variation of a classical infinitely repeated principal-agent game with a long-term action that allows a player to make her unilaterally vulnerable. It illustrates the vulnerability paradox of Pareto-optimal SPE and how it is resolved with positive negotiation probabilities.

In each period, the agent chooses effort $e \in [-x, \bar{e}]$. The principal's stage game payoff is simply given by e. Positive effort level generate value for the principal, negative levels harm her. The state x describes the principal's vulnerability, i.e. the maximum harm that the agent can inflict on the principal. The agent has effort costs $k(e)$ that are strictly increasing and strictly convex for positive effort levels. Zero effort or harming the principal is costless, i.e. $k(e) = 0$ for $e \leq 0$. The joint stage game payoffs $\Pi(e, x) = e - k(e)$ shall be strictly increasing in e for all $e \in [-x, \bar{e}]$. All monetary transfers and effort levels are perfectly observed by the principal and agent, but no enforceable contracts can be written.

The game starts in an initial state $x_0 = 0$, in which the agent cannot harm the principal. In state x_0, the principal has the option to make herself permanently more vulnerable. For example, the principal can grant the agent access to some critical internal infrastructure and grant legal permission that allows the agent to freely damage the principal. If the principal chooses that option, the game moves to a state $x_1 > 0$

in which the principal is vulnerable. For simplicity, the game then stays in state x_1 forever. As long as the principal has not made herself vulnerable, the state remains x_0.

The following straightforward result contains the gist of the vulnerability paradox.

Fact 2. *The principal agent-game always has a Pareto-optimal SPE in which the principal makes herself unilaterally vulnerable in period 1. If the discount factor δ is sufficiently low, such that maximum effort $\bar{e}$ cannot be implemented in the no-vulnerability state x_0, every Pareto-optimal SPE prescribes that the principal makes herself unilaterally vulnerable in period 1.*

Pareto-optimal SPE ignore the possibility that the principal's vulnerability may be exploited by the agent. The players are assumed to always coordinate away from any such exploitation on the equilibrium path. The agent will only harm the principal off-equilibrium path as a punishment if the principal deviates from a prescribed payment. Since tougher punishments allow to implement higher payments and thus higher effort levels, a principal will always make herself more vulnerable in a Pareto-optimal SPE if first-best effort $\bar{e}$ cannot already be implemented in state x_0. In fact, for every discount factor $\delta > 0$ first best effort $\bar{e}$ can be implemented if the vulnerability x_1 is high enough.

We consider the assumption that relational contracts can coordinate away from any attempt to exploit such vulnerabilities as fairly implausible. It seems much more plausible that vulnerable parties worry that they will be held up and have a worse bargaining position in future interactions.

How Repeated Negotiations Resolve The Vulnerability Paradox

Our concept of repeated negotiation equilibrium incorporates such concerns. With positive negotiation probabilities a natural trade-off arises. The principal will only make herself vulnerable if the efficiency gain from the tougher punishment option outweighs the deterioration of her bargaining position in future negotiations. We will formally analyze this trade-off.

Consider a repeated game version of our principal-agent relationship with a fixed state x and a discount factor of $\tilde{\delta} = \delta(1-\rho)$. Let $\tilde{U}(x)$ denote the highest joint SPE payoff of that repeated game. The principal's (expected) negotiation payoffs in such a repeated game are given by

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

On the one hand, a higher vulnerability x, increases the principal's negotiation payoff, since higher joint payoffs $\tilde{U}(x)$ can be implemented. One the other hand, it reduces the principal's negotiation payoffs, since the possibility of stronger harm worsens her bargaining position.

To study RNE in the stochastic game with endogenous vulnerability, it is instructive to first look at the case with a fixed negotiation probability $\rho > 0$ and the limit of no discounting $\delta \to 1$. All continuation payoffs in a pure strategy RNE will then always be approximately equal to the negotiation payoffs of the state, the players end up in the long run. Consequently, the principal will make herself vulnerable if and only if she has higher negotiation probabilities in the vulnerable state:

$$\tilde{r}_1(x_1) \geq \tilde{r}_1(x_0) \Leftrightarrow$$
$$\beta_1 \left(\tilde{U}(x_1) - \tilde{U}(x_0) \right) \geq (1 - \beta_1)(x_1 - x_0)$$

The difference to a Pareto-optimal SPE is not that the principal will never make herself vulnerable, but that she faces a *trade-off* between the efficiency gain and the worsening of her bargaining position. This trade-off is affected by the principal's exogenously specified bargaining weight β_1. If she has a larger bargaining weight, she can appropriate a larger share of the joint surplus and thus puts larger weight on the efficiency gains of vulnerability.

For discount factors strictly below 1, the results stay qualitatively the same but are complicated by the more complex structure of continuation payoffs.

Proposition 3. *For a positive negotiation probability and in the limit $\delta \to 1$, the principal makes herself vulnerable if and only if her negotiation payoff in the repeated game with exogenously given state is larger when vulnerable, i.e. $\tilde{r}_1(x_0) \leq \tilde{r}_1(x_1)$. More generally, the principal makes herself vulnerable in an RNE only if either her bargaining weight β_1 is sufficiently high or the negotiation probability ρ is sufficiently low.*

This example can be interpreted as a repeated game augmented by an endogenous decision about institutional design. Here the design choice just consists of a level of vulnerability, but in more complex examples there can be many institutional factors that affect bargaining positions. The key contribution of RNE compared to Pareto-optimal SPE is that in such situations parties will care about how the selected institutions affect their future bargaining positions. This, in turn, will affect the shape of the institutions that endogenously arise.

4.2 Mutual vs Unilateral Vulnerability

We now consider a variant of the previous example in which players can make themselves unilaterally or mutually vulnerable. We consider a symmetric stage game with two players that have equal bargaining weights $\beta_1 = \beta_2 = \frac{1}{2}$. Each player i chooses an effort level $e_i \in [-d_j, \bar{e}]$ where $d_j > 0$ denotes the maximum damage that can be inflicted on the other player j. Payoffs are $\pi_i = e_j - k(e_i)$. As in the previous example, the effort cost function $k(e_i)$ is strictly continuously increasing and strictly convex for positive effort levels, while zero effort or damage infliction is costless. Joint profits shall be strictly increasing in e_1 and e_2.

There are four different states $\{x_0, x_1, x_2, x_{12}\}$. In state x_0 no damage can be inflicted $d_1 = d_2 = 0$. In state x_1 only player 1 is vulnerable: $d_1 = \bar{d}$ and $d_2 = 0$ with $\bar{d} > 0$. In state x_2 only player 2 is vulnerable: $d_1 = 0$ and $d_2 = \bar{d}$. In state x_{12} both players are vulnerable: $d_1 = d_2 = \bar{d}$.

In state x_0 each player can decide to do nothing, make herself unilaterally vulnerable, or propose a mutual vulnerability. If a single player i makes herself unilaterally vulnerable, the game moves to state x_i and stays there forever. If both players propose mutual vulnerability (or both make themselves simultaneously unilaterally vulnerable) the state moves to x_{12} and stays there forever. Otherwise the state remains x_0. This means if only one player proposes to implement mutual vulnerability, while the other player does nothing, no player becomes vulnerable and the state remains x_0. Creating a unilateral vulnerability is costless. Implementing mutual vulnerability shall involve a small cost $\varepsilon > 0$ for each player that reflects coordination and monitoring effort to ensure that indeed both players simultaneously implement the vulnerability.

We assume that the maximum damage $\bar{d}$ is high enough, such that unilateral vulnerability suffices to implement first best efforts $e_1 = e_2 = \bar{e}$ in a SPE given a discount factor of $\tilde{\delta} = (1-\delta)\rho$. Yet, without any vulnerabilities (state x_0), first best efforts shall not be implementable, even under the discount factor of δ. From these assumptions follows directly

Fact 3. *In every Pareto-optimal SPE at least one player makes herself unilaterally vulnerable in period 1.*

In a symmetric Pareto-optimal SPE both players make themselves simultaneously unilaterally vulnerable in period 1. However, they are not willing to incur a small cost ε to guarantee that they only become vulnerable if also the other player becomes vulnerable. The reason is, like in the previous example, that in a Pareto-optimal SPE there is no harm to be unilaterally vulnerable since players are assumed to coordinate away from any exploitation of the vulnerability on the equilibrium path.

In contrast, we find with positive negotiation probabilities

Proposition 4. *If the negotiation probability ρ is positive, implementation cost ε are sufficiently small and the discount factor δ and maximum damage $\bar{d}$ are sufficiently large, both players will propose and implement mutual vulnerability in period 1 in an RNE.*

With positive negotiation probabilities, a vulnerable player suffers from a weak future bargaining position while the other player benefits. If future negotiations have sufficient weight, a simple promise that both players make themselves simultaneously vulnerable is therefore not credible. Consequently, in an RNE players are willing to incur monitoring and implementation costs to ensure that vulnerabilities are indeed mutually implemented.

It is also instructive to look at a variant of the example, in which there is no external mechanism to ensure mutual implementation of vulnerability, but each party can choose each period a continuous value Δd_i of how much they want to increase their current level of vulnerability. Pareto-optimal SPE will prescribe that players immediately increase their vulnerability up to the levels required to implement first best efforts. In contrast, in an RNE with negotiation probability strictly between 0 and 1, parties will gradually and symmetrically increase their vulnerabilities in small steps each period. Too large steps will not be incentive compatible given the prospect of future negotiations. Yet, small increases in vulnerabilities remain incentive compatible, since deviations yield only small gains in future negotiations but can be punished during the time until the next negotiation takes place.

Kopányi-Peuker, Offerman and Sloof (2017) show in an experiment that if players can choose how much they can be punished for not cooperating in a subsequent prisoners' dilemma, then a gradual mechanism with incremental and conditional increases in vulnerability is more effective than simple simultaneous choices of the vulnerability level. In applications, higher vulnerability often is a consequence of tighter integration, e.g. when manufacturers and upstream suppliers build up a just-in-time supply chain that reduces the amount of inventory. Repeated negotiations provide one reason why in relationships such tighter integration will gradually evolve over time.[14]

[14]Another reason to start small in relationships is that parties need time to learn each other's types

An example of the use of costly monitoring and compliance devices to ensure mutual vulnerability is the Anti-Balistic Missiles Treaty from 1972-2002 between the US and the Soviet Union. By limiting the number of anti-balistic missiles, parties ensured a state of mutual vulnerability with respect to each others nuclear strikes. Platt (1991) explains how the parties took considerable effort to prevent any deviation from this state of mutual vulnerability. Thirteen of the treaty's sixteen articles were intended to prevent any deviation. Sophisticated and costly monitoring devices on land, sea, and in space were deployed to monitor compliance with the treaty.

4.3 A simple arms race

This example shows how future negotiations create incentives for costly investments that have the sole goal of improving a party's bargaining position and no other direct benefit. It also illustrates that equilibrium payoffs do not necessarily weakly decrease in the negotiation probability, as is always the case in a repeated game.

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 acquisition is successful may be determined stochastically. Assume that an attempt is successful with probability $\phi > 0$ and that unsuccessful attempts to acquire a weapon are not observed. Once a weapon has been successfully acquired, play moves from state x_0 to state x_1, in which the player who acquired a weapon can use it at some cost $c > 0$ to inflict a damage $d > 0$ on the other player. There are no direct benefits from using a weapon. Assume for simplicity that only player 1 can acquire a weapon and that the weapon can be used as often as desired. If no weapons are bought or used, players get a payoff of zero.

Proposition 5. *In the unique Markov perfect equilibrium outcome, as well as in all Pareto-optimal PPE 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.*

That weapons are not build in the two extreme cases of no repeated negotiation or negotiation in every period has different reasons. It is evident that costly acquisition

in a framework of incomplete information, see Watson (1999) and Watson (2002).

of weapons on the equilibrium path cannot be part of a Pareto-optimal PPE, 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 the benefits of having weapons in future negotiations are simply too high. Basically, with an intermediate negotiation probability, an initial resolve to ignore threats of weapon usage unless money is paid eventually fades.

The example illustrates that to prevent an arms race, it is not sufficient that there are no direct gains from using weapons. The induced socially inefficient investments can be interpreted as a particular incarnation of a hold-up problem. Note that if both players can acquire a weapon (i.e. play can also move from state x_1 in which only one player has a weapon to a state x_2 in which both player possess one) then this may eventually reduce negotiation payoffs (e.g. to zero with equal bargaining power) and therefore also reduce incentives to acquire a weapon in the first place.

4.4 Inside options versus outside options

An important insight of non-cooperative bargaining models 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. 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, Shaked and Sutton (1989)). The difference between outside options and inside options can have important implications for hold-up problems and optimal asset ownership (de Meza and Lockwood (1998)). In contrast, these differences typically do not matter for traditional models of relational contracting.

To see how repeated negotiations naturally extend the outside option principle to relational contracting, consider a variation of the repeated principal-agent game. Within the relationship, the principal's and agent's stage game payoffs shall be given by

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

The parameter $e \in [0, \bar{e}]$ 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 effort is chosen.

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 each player i gets in the current and all future periods an outside option payoff of π_i^{oo}. We assume that $\pi_i^{io} < \pi_i^{oo}$, 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 probability $\varepsilon > 0$ that the break up is not successful and players remain in the relationship next period.

Proposition 6. *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*

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

Given a positive negotiation probability $\rho > 0$, then in the limit $\delta \to 1$, expected negotiation payoffs satisfy instead the outside option principle: Unless the outside option is binding for some player i, they are solely determined by the inside options and given by

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

These results are straightforward. The lowest subgame perfect equilibrium payoff of each player is given by her outside option. In the absence of repeated negotiations, 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 $\delta \to 1$, 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. In this case, also the maximum implementable joint payoff $\bar{U}$ only depends on the inside options, not on the outside options.

Thus, if one would augment this game for long-term, institutional-design actions that can influence both inside and outside options, repeated negotiations emphasize the importance of the effect of institutions on the inside options compared to the effect on outside options.

4.5 The blackmailing game

We now consider a simple game that illustrates why repeated negotiation equilibria can have irregular 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

$$\begin{aligned} \pi(a_S, x_0) &= (0, 1) \\ \pi(x_1) = \pi(a_R, x_0) &= (0, 0). \end{aligned} \tag{8}$$

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

$$\begin{aligned} \tilde{r}_1(x_0) &= \beta_1 \\ \tilde{r}_2(x_0) &= 1 - \beta_1. \end{aligned} \tag{9}$$

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 an irregular expected negotiation payoff of zero in state x_0. 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 $\Gamma(r, x_0)$ 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 negotiation outcomes only depend on the state, both players know that when negotiation 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.

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 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.

Brinkmanship The blackmailer could extort larger payments if the game allows to conduct brinkmanship (Schelling (1960) and Schwarz and Sonin (2008)).[15] The blackmailer needs an observable action that reveals the evidence with positive probability smaller than 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 inform the target about it. There is a positive probability that the envelope will still be lying on the street if the target comes to fetch it, but the envelope might already have been put into the postal box by some helpful minded pedestrian. Hence, in the following we assume that there is an observable, pure action a_ϕ that reveals the evidence only with probability $\phi \in (0,1)$.

Proposition 7. *For all $\phi \in (0,1)$, there exists an RNE in which the brinkmanship action a_ϕ is used to extort a payment from the target. This payment is maximized if ϕ is equal to $\phi^* = \frac{(1-\delta)(1-\rho)}{\rho\beta_1+(1-\delta)(1-\rho)}$, with the blackmailer's negotiation payoff equal to the regular negotiation payoff $\beta_1 \frac{(1-\rho)(1-\delta(1-\rho))}{\rho(\beta_1-1)+1-\delta(1-\rho)}$.*

A larger value of ϕ means a harsher punishment which can be used to extract larger payments. If ϕ is relatively small, the blackmailer's negotiation payoff increases in ϕ. However, if ϕ becomes too large, credibility of the punishment imposes an upper bound on the negotiation payoff. This explains the result that an intermediate value will be optimal. As the blackmailer's bargaining power β_1 increases, his commitment problem gets more severe and the optimal brinkmanship action has a lower probability of revealing the evidence, i.e., optimal punishment becomes more gradual. The same holds for as an increasing frequency of renegotiation ρ or a larger weight on future payoffs δ. Conversely, as the negotiation probability ρ converges to zero, the optimal punishment probability converges to 1 and the corresponding blackmailer's negotiation payoff converges to β_1.

[15] Brinkmanship is the ability of an aggressor to choose an observable action that leads with some positive probability to a mutually undesirable outcome. Schwarz and Sonin (2007) show that such a divisible threat can dramatically increase the bargaining value of a non-credible threat by making punishment possible in a subgame-perfect equilibrium. In contrast, in our example there exists SPE in which the the evidence is revealed. A commitment problem instead results from positive negotiation probabilities.

5 Concluding Remarks

This paper has been motivated by the discrepancy between the behavioral assumptions of traditional relational contracting models and hold-up models. One already known facet of this discrepancy is that Pareto-efficient SPE render most hold-up problems void. More importantly, we have exemplified with several infinite horizon games that have long-term actions how Pareto-efficient SPE yield implausible results, which can be resolved by adding a "hold-up component" to relational contracts.

Our formulation of this hold-up component assumes that with an exogenously given probability, an existing relational contract is newly negotiated. The notion that relational contracts are randomly newly negotiated facilitates interpretation and allows a definition applicable to a general class of games. Yet, randomization is not essential in the sense that all insights only rely on the impact of future negotiations on *expected* continuation payoffs. The key feature is that players care about the effects of long-term actions on their future bargaining positions. We have illustrated how this feature yields more plausible results than Pareto-optimal PPE, e.g. it resolves the vulnerability paradox, can explain costly arms races even if there is no direct benefit from using arms, and reestablishes the outside-option principle in repeated principal-agent interactions.

The formalization of negotiations as a random-dictator bargaining is due to the technical challenge to develop an equilibrium concept that is tractable for and exists in general discounted stochastic games with transfers. An interesting area of future research is the exploration of alternative bargaining formulations. One possible route is to formulate a theory of endogenous disagreement points, as has been done for repeated games by Miller and Watson (2013). Yet, there are considerable challenges concerning tractability and existence for such alternative formulations in an environment of discounted stochastic games.

A straightforward extension of our model are to allow negotiation probabilities $\rho(x)$ that depend on the current state. Our existence result also holds for this more general model. One avenue of future research is to explore models with state-dependent negotiation probabilities. One example is to model incompleteness of relational contracts. To model the fact that some state is not considered in an initial relational contract, one can assign a negotiation probability of 1 to such states while having smaller or zero negotiation probabilities for states that have been initially considered. State-dependent negotiation probabilities can also allow for models with endogenously arising negotiations, in which players have to actively move to states in which negotiations can take place. Other applications could entail models in which negotiation probabilities decrease over time or after streaks of successful cooperation. This can formulate the idea that longer relationships become more stable over time.

Bibliography

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti, "Renegotiation and Symmetry in Repeated Games," *Journal of Economic Theory*, 1993, *60* (2), 217–240.

Aghion, Philippe, Mathias Dewatripont, and Patrick Rey, "Renegotiation Design with Unverifiable Information," *Econometrica*, March 1994, *62* (2), 257–282.

Asheim, Geir B, "Extending renegotiation-proofness to infinite horizon games," *Games and Economic Behavior*, August 1991, *3* (3), 278–294.

Avery, Christopher and Peter B. Zemsky, "Money Burning and Multiple Equilibria in Bargaining," *Games and Economic Behavior*, September 1994, *7* (2), 154–168.

Baker, George, Robert Gibbons, and Kevin J. Murphy, "Subjective Performance Measures in Optimal Incentive Contracts," *The Quarterly Journal of Economics*, 1994, *109* (4), 1125–1156.

——, ——, and ——, "Relational Contracts and the Theory of the Firm," *Quarterly Journal of Economics*, February 2002, *117* (1), 39–84.

Bernheim, Douglas B. and Debraj Ray, "Collective dynamic consistency in repeated games," *Games and Economic Behavior*, December 1989, *1* (4), 295–326.

Binmore, Ken, Avner Shaked, and John Sutton, "An Outside Option Experiment," *Q J Econ*, November 1989, *104* (4), 753–770.

Blonski, Matthias and Giancarlo Spagnolo, "Relational Contracts and Property Rights," *working paper*, 2007.

Board, Simon, "Relational Contracts and the Value of Loyalty," *The American Economic Review*, 2011, *101* (7), 3349–3367.

Bull, Clive, "The Existence of Self-Enforcing Implicit Contracts," 1987, *102* (1), 147–159.

Chatterjee, K. and L. Samuelson, "Perfect equilibria in simultaneous-offers bargaining," *Int J Game Theory*, September 1990, *19* (3), 237–267.

Che, Yeon-Koo and Donald B. Hausch, "Cooperative Investments and the Value of Contracting," *The American Economic Review*, 1999, *89* (1), 125–147.

—— **and József Sákovics**, "A Dynamic Theory of Holdup," *Econometrica*, 2004, *72* (4), 1063–1103.

de Meza, David and Ben Lockwood, "Does Asset Ownership Always Motivate Managers? Outside Options and the Property Rights Theory of the Firm," *Q J Econ*, May 1998, *113* (2), 361–386.

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

Ellingsen, Tore and Jack Robles, "Does Evolution Solve the Hold-Up Problem?," *Games and Economic Behavior*, April 2002, *39* (1), 28–53.

—— **and Magnus Johannesson**, "Is There a Hold-up Problem?," *Scandinavian Journal of Economics*, 2004, *106* (3), 475–494.

Evans, Robert, "Simple Efficient Contracts in Complex Environments," *Econometrica*, 2008, *76* (3), 459–491.

Farrell, Joseph and Eric Maskin, "Renegotiation in repeated games," *Games and Economic Behavior*, 1989, *1* (4), 327–360.

Garvey, Gerald T., "Why reputation favors joint ventures over vertical and horizontal integration A simple model," *Journal of Economic Behavior & Organization*, December 1995, *28* (3), 387–397.

Goldlücke, Susanne and Sebastian Kranz, "Infinitely repeated games with public monitoring and monetary transfers," *Journal of Economic Theory*, 2012, *147* (3), 1191–1221.

—— **and** —— , "Discounted stochastic games with voluntary transfers," *Econ Theory*, June 2017, pp. 1–29.

Grout, Paul A., "Investment and Wages in the Absence of Binding Contracts: A Nash Bargaining Approach," *Econometrica*, 1984, *52* (2), 449–460.

Halac, Marina, "Investing in a relationship," *The RAND Journal of Economics*, 2015, *46* (1), 165–185.

Halonen, Maija, "Reputation and the Allocation of Ownership*," *The Economic Journal*, July 2002, *112* (481), 539–558.

Hart, Oliver and John Moore, "Incomplete Contracts and Renegotiation," *Econometrica*, 1988, *56* (4), 755–785.

Iossa, Elisabetta and Giancarlo Spagnolo, "Contracts as Threats: On a Rationale for Rewarding A while Hoping for B," *CEPR Discussion Papers DP8195*, 2011.

Klein, Benjamin, Robert G. Crawford, and Armen A. Alchian, "Vertical Integration, Appropriable Rents, and the Competitive Contracting Process," 1978, *21* (2), 297–326.

Kopányi-Peuker, Anita, Theo Offerman, and Randolph Sloof, "Fostering cooperation through the enhancement of own vulnerability," *Games and Economic Behavior*, 2017, *101* (Supplement C), 273 – 290. Special Issue in Honor of John O. Ledyard.

Levin, Jonathan, "Multilateral Contracting and the Employment Relationship," *Q J Econ*, August 2002, *117* (3), 1075–1103.

____, "Relational Incentive Contracts," *American Economic Review*, June 2003, *93* (3), 835–857.

MacLeod, W. Bentley and James M. Malcomson, "Implicit Contracts, Incentive Compatibility, and Involuntary Unemployment," *Econometrica*, 1989, *57* (2), 447–480.

Malcomson, James M., "Relational Incentive Contracts," in ROBERT GIBBONS and JOHN ROBERTS, eds., *The Handbook of Organizational Economics*, Princeton University Press, 2012, pp. 1014–1065.

Miller, David A. and Joel Watson, "A Theory of Disagreement in Repeated Games With Bargaining," *Econometrica*, November 2013, *81* (6), 2303–2350.

Nöldeke, Georg and Klaus M. Schmidt, "Option Contracts and Renegotiation: A Solution to the Hold-up Problem," *The RAND Journal of Economics*, 1995, *26* (2), 163–179.

Pearce, David G and Ennio Stacchetti, "The Interaction of Implicit and Explicit Contracts in Repeated Agency," *Games and Economic Behavior*, April 1998, *23* (1), 75–96.

Pitchford, Rohan and Christopher M. Snyder, "A solution to the hold-up problem involving gradual investment," *Journal of Economic Theory*, January 2004, *114* (1), 88–103.

Platt, Alan, *The Anti-Ballistic Missile Treaty*, New York: Palgrave Macmillan US,

Ramey, Garey and Joel Watson, "Contractual Fragility, Job Destruction, and Business Cycles," *Q J Econ*, August 1997, *112* (3), 873–911.

Rubinstein, Ariel, "Perfect Equilibrium in a Bargaining Model," *Econometrica*, 1982, *50* (1), 97–109.

Safronov, Mikhail and Bruno Strulovici, "Strategic Renegotiation in Repeated Games," 2016.

Schelling, Thomas C., *The Strategy of Conflict*, Harvard University Press, 1960. Google-Books-ID: 7RkL4Z8Yg5AC.

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

Schwarz, Michael and Konstantin Sonin, "A Theory of Brinkmanship, Conflicts, and Commitments," *J Law Econ Organ*, May 2008, *24* (1), 163–183.

Sobel, Matthew J., "Noncooperative Stochastic Games," *Ann. Math. Statist.*, December 1971, *42* (6), 1930–1935.

Tian, Guoqiang and Jianxin Zhou, "The maximum Theorem and the existence of Nash equilibrium of (generalized) games without lower semicontinuities," *Journal of Mathematical Analysis and Applications*, 1992, *166* (2), 351 – 364.

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

Watson, Joel, “Starting Small and Renegotiation,” *Journal of Economic Theory*, March 1999, *85* (1), 52–90.

——, “Starting Small and Commitment,” *Games and Economic Behavior*, 2002, *38* (1), 176 – 199.

Williamson, Oliver E., *The Economic Institutions of Capitalism*, Simon and Schuster, 1985. Google-Books-ID: MUPVLuiy9uQC.

Appendix A: Proofs

Proof of Lemma 1

Let $\mathcal{R}$ denote the space of all negotiation payoffs. 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.

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

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 r and $\tilde{r}$ satisfy $r_j(x) \leq \tilde{r}_j(x)$ for all j, x then $g_j(x|r, \sigma) \leq g_j(x|\tilde{r}, \sigma)$ for all 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 j, x

$$g_j(x|\hat{r} + K, \sigma) \leq g_j(x|\hat{r}, \sigma) + \gamma K$$

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 $\gamma = \delta$.

Proof of Lemma 2

To prove this result, we first prove that for any incentive compatible contract profile σ (meaning that all $\sigma_{(i,x)}$ are PPE in $\Gamma(r(.|\sigma), x)$), there exists an incentive compatible simple contract profile $\tilde{\boldsymbol{\sigma}}$ which has the same negotiation payoffs $r(.|\sigma) = r(.|\tilde{\sigma})$.

We know from Goldlücke and Kranz (2017) that for every state x there exists an optimal simple equilibrium $\tilde{\sigma}_{(i,x)}$ that implements the same payoffs as $\sigma_{(i,x)}$ in the truncated game $\Gamma(r(.|\sigma), x)$. Note that the action plan is independent of the subscript i. A profile $\tilde{\sigma}$ of such optimal simple renegotiation outcomes satisfies

$$g(x|r(.|\sigma), \sigma) = g(x|r(.|\sigma), \tilde{\sigma}) \text{ for all } x \in X$$

It follows from Lemma 1 that two contract profiles σ and $\tilde{\sigma}$ induce the same negotiation payoffs if and only if they implement the same payoffs in the truncated games given $r(.|\sigma)$, i.e., $g(.|r(.|\sigma), \sigma) = g(.|r(.|\sigma), \tilde{\sigma})$. Hence this directly implies

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

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

We can now show the result of the lemma by proving the following claim: 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 maximum of her continuation payoffs across all relational contracts that are incentive compatible given $\sigma_{-(i,x)}$.

Let $\bar{r}_i^i$ denote the supremum of player i's payoffs in state x that can be implemented with incentive compatible relational contracts. Since every incentive compatible relational contract chosen by player i in state x can be replaced by an optimal simple relational contract that has the same negotiation payoffs, there exists a sequence $\{\tilde{\sigma}_{(i,x)}^m\}_{m=1}^{\infty}$ of incentive compatible simple optimal relational contracts such that the corresponding payoffs for player i converge to $\bar{r}_i^i$. 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 to verify that $\bar{\sigma}_{(i,x)}$ is an optimal simple equilibrium in the truncated game $\Gamma(r(\bar{\sigma}_{(i,x)}, \sigma_{-(i,x)}), x)$.

Proof of Theorem 1

We first prove the following lemma.

Lemma 3. *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, importantly, also 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. $\square$

We now prove the existence theorem. 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. The players in these auxiliary games are the combinations (i,x) of a player and state of the original stochastic game. Their action space is the set of simple relational contracts. More precisely, players choose up-front transfers p^0, a (possibly) mixed action profile $\alpha^k(x)$ for each state x and phase $k = e, 1, ..., n$, and a payment plan $p^k(x, y, x')$ for each combination of signals y and states x, x' such that x' can be reached from x, with the additional requirement that $\sum_{i=1}^n p_i^k(x, y, x') \geq 0$ for all x, y, x', k. With appropriately chosen bounds on payments this strategy set is a compact and convex set. To define a player's payoff $\pi_{i,x}(\sigma)$ for a given strategy profile σ, we first describe in detail what it means that the simple relational contract $\sigma_{i,x}$ is incentive compatible given $\sigma_{-i.x}$. If we add up the gains from deviation for any player, current stage (payment or action stage), phase $k \in \{e, 1, ..., n\}$, the previously realized signal y, the current state x', the previous state x'', and a pair of player and state (i, x) that identifies the current relational contract, then this number has to be zero for the relational contract profile σ. Concretely, we define for given σ the gain from a deviation from a required payment in state x' as $\Delta^p_{i,x} = u^j_j(x') - u^e_j(x') + (1-\delta)p^k_j(x'', y, x')$ and the maximum gain as

$$\Delta^p_{i,x}(\sigma) = \max_{k \in \{e,1,...,n\}} \max_{j \in \{1,..,n\}} \max_{y \in Y} \max_{x', x'' \in X} \Delta^p(k, j, y, x', x'', \sigma).$$

Similarly, the maximum gain from deviations in the action stage is

$$\Delta^a_{i,x}(\sigma) = \max_{k \in \{e,1,...,n\}} \max_{j \in \{1,..,n\}} \max_{x' \in X} \max_{a'_j \in A_j(x')} \Delta^a(k, j, a^{'}_j, x', \sigma).$$

where $\Delta^a_{i,x} = (1-\delta)\pi_j(x', a'_j, \alpha^k_{-j}) + \delta E[(1-\rho)(u^e_j - (1-\delta)p^k_j) + \rho r_j | x', a^{'}_j, a^k_{-j}] - u^k_j(x')$. It then holds that$\sigma_{i,x}$ is a PPE in $\Gamma(r(.|\sigma), x)$ if and only if

$$\Delta_{i.x}(\sigma) = \max\{\Delta^p_{i,x}(\sigma), \Delta^a_{i,x}(\sigma)\} = 0.$$

One can show that $\Delta_{i,x}$ is a quasi-convex and continuous function. We now define the payoff function of the auxiliary game as

$$\pi^m_{i,x}(\sigma) = \min\{u_i(\sigma|x, i), K - m\Delta_{i,x}(\sigma)\},$$

which is a quasi-concave and continuous function. 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 at least one Nash equilibrium. For each m, select one of these equilibria, denoted by σ^m. As a 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^*. Because of Lemma 3, each player has a simple relational contract $\sigma_{i,x}$ that is incentive compatible given $\sigma^*_{-i,x}$ and $\sigma^m_{-i,x}$. Since violations of incentive constraints are exceedingly costly as $m \to \infty$, it thus follows from our construction that $\Delta_{i,x}(\sigma^m) \to 0$ as $m \to \infty$, i.e., $\sigma^*_{i,x}$ must be incentive compatible given $\sigma^*_{-i,x}$. Moreover, for any i, x and $\tilde{\sigma}_{i,x}$ that is stable incentive compatible given $\sigma^*_{-i,x}$, there must exist a subsequence

$\sigma^{m_k}_{-i,x} \to \sigma^*_{-i,x}$ as well as $\tilde{\sigma}^k_{i,x} \to \tilde{\sigma}_{i,x}$, such that $\tilde{\sigma}^k_{i,x}$ is a PPE in $\Gamma(r(.|\tilde{\sigma}^k_{i,x}, \sigma^{m_k}_{-i,x}), x)$. Since σ^{m_k} is a Nash equilibrium in the game with payoff functions π^{m_k}, it must be true that $u_i(\sigma^{m_k}|x,i) \geq u_i(\tilde{\sigma}^k_{i,x}, \sigma^{m_k}_{-i,x}|x,i)$ and hence in the limit as $k \to \infty$ this implies $u_i(\sigma^*|x,i) \geq u_i(\tilde{\sigma}_{i,x}, \sigma^*_{-i,x}|x,i)$.

To prove that there exists a canonical RNE, we have to study another artificial game, one with two players. One of them (player 2) takes the role of all the players (i,x) as in the first part of the proof, and the other (player 1) tries to match the negotiation payoffs with a canonical contract profile. We will write this as a generalized game. Player 2's strategy set is the set of all simple relational contract profiles and his payoff function is equal to

$$\pi^m_2(\tilde{\sigma}, \sigma) = \sum_{i,x} \pi^m_{i,x}(\sigma_{i,x}, \tilde{\sigma}_{-i,x}).$$

That is, for each (i,x) the player takes as given the choice of player 1. Player 1's choice will be a simple relational contract profile $\tilde{\sigma}$ with the propery that for all states x, the relational contracts $\tilde{\sigma}_{i,x}$, $i = 1,..,n$, only differ by their up-front payments. We denote the (compact and convex) set of such relational contract profiles by $\mathcal{C}$. Let $U^0(\sigma|x,i)$ denote the joint payoffs on the equilibrium path if players follow the strategy profile $\sigma \in \mathcal{C}$ after they make the up-front payments in state x as specified by σ, i.e., these are joint payoffs following history $x, i, p^0(x,i)$ with $p^0(x,i) = \sigma(x,i)$. Let $v_j(\sigma|x,i)$ be player j's punishment payoff in state x, i.e. this is player $j's$ payoff if j deviates from the payment specified by σ in state x. Let $V(\sigma|x,i) = \sum_{j=1}^n v_j(\sigma|x,i)$. Player 1 maximizes the function

$$\pi_1(\tilde{\sigma}, \sigma) = \sum_{x,i} (U^0(\tilde{\sigma}|x,i) - V(\tilde{\sigma}|x,i))$$

over all $\tilde{\sigma} \in \gamma(\sigma)$, where the feasible strategy set $\gamma(\sigma)$ is given by

$$\gamma(\sigma) = \{\tilde{\sigma} \in \mathcal{C} \;|\Delta_{i,x}(\hat{\sigma}) \leq \Delta_{i,x}(\sigma), u(\sigma|x,i) = u(\tilde{\sigma}|x,i) \text{ for all } x \in X, i = 1, ..., n\}.$$

For given σ, an element of $\gamma(\sigma)$ can be constructed analogous to the construction of a PPE with the same negotiation payoffs as a given RNE in the proof of Lemma 2.[16] This

[16] More precisely, for each state x there must be a continuation strategy profile of σ which gives the largest joint payoff for players in state x and similarly, there must be continuation strategy profiles that give the lowest payoffs for each player. Conequently, there must be action profiles $(\tilde{\alpha}^e(x))_{x \in X}$ and continuation payoffs $w^e(x,y,x')$ such that $U^e(x)$ is maximized and action profiles $(\tilde{\alpha}^i(x))_x$ and continuation payoffs $w^i(x,y,x')$ such that $v_i(x)$ is minimized. These $\tilde{\alpha}^k$ define the action plan of $\tilde{\sigma}$.To make sure that the maximum gain from deviation in the newly constructed $\tilde{\sigma}$ is not larger than in the original σ, we define $(1-\delta)\tilde{p}^k(x,y,x') = u^e_i(x') - w^k_i(x,y,x')$, which equalizes continuation payoffs

means that $\gamma(\sigma)$ is nonempty. Because u is linear and $\Delta_{i,x}$ is quasi-convex, the set $\gamma(\sigma)$ is convex, and because $\Delta_{i,x}$ and u are continuous, it is compact. Since u and $\Delta_{i,x}$ are continuous, the correspondence γ is continuous. Moreover, $\sum_{x,i}(U^0(\tilde{\sigma}|x,i) - V(\tilde{\sigma}|x,i))$ is continous and linear in $\tilde{\sigma}$. Applying an equilibrium existence result for generalized games (e.g. Tian and Zhou (1992)) for each $m \in \mathbb{N}$, we can define a sequence of equilibria of the generalized game $(\tilde{\sigma}^m, \sigma^m)_{m=1}^{\infty}$ which must have a limit (of a subsequence) that we denote by $(\tilde{\sigma}^*, \sigma^*)$, with $\tilde{\sigma}^* \in \mathcal{C}$. For the fixed points in the sequence it must hold that$\Delta_{i,x}(\tilde{\sigma}^m) \leq \Delta_{i,x}(\sigma^m)$ and $u(\sigma^m|x,i) = u(\tilde{\sigma}^m|x,i)$, and because $\Delta_{i,x}(\sigma^m) \to 0$ then also $\Delta_{i,x}(\tilde{\sigma}^m) \to 0$, which means that $\tilde{\sigma}^*$ must be a PPE in the original game. Moreover, for any i, x and $\hat{\sigma}_{i,x}$that is strongly incentive compatible given $\tilde{\sigma}_{-i,x}$we can infer with similar steps as above that $u_i(\sigma^*_{i,x}, \tilde{\sigma}^*_{-i,x}|x,i) \geq u_i(\hat{\sigma}_{i,x}, \tilde{\sigma}^*_{-i,x}|x,i)$ which also means that $u_i(\tilde{\sigma}^*|x,i) \geq u_i(\hat{\sigma}_{i,x}, \tilde{\sigma}^*_{-i,x}|x,i)$. It remains to show that the simple equilibrium $\tilde{\sigma}^*_{i,x}$ is optimal in the truncated game $\Gamma(r(.|\tilde{\sigma}^*), x)$, but this holds because all $\tilde{\sigma}^m$ maximize $\sum_{x,i}(U^0(\hat{\sigma}|x,i) - V^0(\hat{\sigma}|x,i))$ among all those with the same negotiation payoff as σ^m and with sufficiently low deviations from the equilibrium conditions.

Proof of Proposition 1

Consider some state $x \in X$ in a stochastic game with monotone state transitions. 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 equilibrium of the truncated game $\Gamma(r,x)$ and implements the original negotiation payoff r without violating the RNE conditions.

Proof of Proposition 2

For a given strategy profile σ, let $\tilde{\pi}(t,\sigma)$ denote the expected payoffs in period t in a repeated game with zero negotiation probability. The expected payoff in the truncated

following play of $\tilde{\alpha}^k$ in σ and $\tilde{\sigma}$. In particular, if σ is a PPE, then also the constructed $\tilde{\sigma}$ would be a PPE. To finish the construction of $\tilde{\sigma}$, we define $\tilde{p}^0(x,i)$ such that $u(\sigma|x,i) = u(\tilde{\sigma}|x,i)$ for all $x \in X, i = 1, ..., n$.

game $\Gamma(r,x)$ can be written as

$$u(\sigma) \quad = \quad (1-\delta)\sum_{t=0}^{\infty}(\delta(1-\rho))^t\,\tilde{\pi}(t,\sigma)+r\frac{\delta\rho}{1-\delta+\delta\rho}. \tag{10}$$

This is a positive affine transformation of a discounted payoff of the repeated game with discount factor $\tilde{\delta}=(1-\rho)\delta$. Hence, 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. That is, each player i will suggest a relational contract that yields $\bar{U}(\tilde{\delta})-\sum_{j\neq i}\bar{v}_j(\tilde{\delta})$ for player i and $\bar{v}_j(\tilde{\delta})$ for each player $j\neq i$ as long as there are no new negotiations. It follows that player i's punishment payoffs satisfy

$$v_i \quad = \quad \frac{1-\delta}{1-\tilde{\delta}}\bar{v}_i(\tilde{\delta})+\frac{\delta-\tilde{\delta}}{1-\tilde{\delta}}r_i \tag{11}$$

Negotiation payoffs must satisfy

$$r_i=v_i+\beta_i(\bar{U}(\tilde{\delta})-\sum_{j=1}^{n}v_j) \tag{12}$$

Combining these two equalities yields

$$v_i=\bar{v}_i(\tilde{\delta})+\frac{\delta-\tilde{\delta}}{1-\tilde{\delta}}\beta_i(\bar{U}(\tilde{\delta})-\sum_{j=1}^{n}\bar{v}_j(\tilde{\delta})) \tag{13}$$

and the expression for r_i in equation (4).

Proof of Proposition 3

This proof and the following proofs use the characterization results for simple equilibria in stochastic games by Goldlücke and Kranz (2017). See Appendix B for a summary. The result for $\delta\to 1$ is shown in the text.

Once state x_1 is reached, players face a repeated game with negotiation payoffs $\tilde{r}(x_1)$ and maximal joint payoffs $\tilde{U}(x_1)$. We consider simple equilibria in the truncated game in which in the equilibrium phase in state x_0, the principal makes herself vulnerable and effort $e_0\geq 0$ is chosen in state x_0. We consider the case that the principal has a higher negotiation payoff in state x_0 than x_1, i.e. $r_1(x_0)\geq\tilde{r}_1(x_1)$, since otherwise there is no problem to incentivize the principal to make herself vulnerable. The joint incentive constraint (SUM-IC) from Proposition 8 then becomes

$$(1-\delta)\left(e_0 - k(e_0)\right) + \delta\tilde{U}(x_1) \geq (1-\delta)e_0 + \delta\rho r_1(x_0) + \delta\rho\tilde{r}_2(x_1)$$

and simplifies to

$$\delta(1-\rho)\tilde{U}(x_1) - (1-\delta)k(e_0) \geq \delta\rho\left(r_1(x_0) - \tilde{r}_1(x_1)\right).$$

A lower bound of $r_1(x_0)$ in any RNE is $\tilde{r}_1(x_0)$, the principal's negotiation payoff in the repeated game with fixed state x_0. That is because the principal can always force to remain in state x_0. By setting $e_0 = 0$, we thus find the following necessary condition for incentive compatibility

$$(1-\rho)\tilde{U}(x_1) \geq \rho\left(\tilde{r}_1(x_0) - \tilde{r}_1(x_1)\right).$$

Using the fact that $\tilde{r}_1(x) = -\left(1-\beta_1\right)x + \beta_1\tilde{U}(x)$, we can rewrite this necessary condition as

$$\tilde{U}(x_1) \geq \frac{\rho}{1-\rho}\Big((1-\beta_1)(x_1 - x_0) - \beta_1(\tilde{U}(x_1) - \tilde{U}(x_0)\Big).$$

Hence, if the principal's bargaining weight is strictly below the threshold

$$\bar{\beta}_1 \equiv \frac{x_1 - x_0}{x_1 - x_0 + \tilde{U}(x_1) - \tilde{U}(x_0)}$$

this condition is violated for sufficiently large negotiation probabilities.

Proof of Proposition 4

If continuation play stays forever in a fixed state in which players have vulnerabilities $d = (d_1, d_2)$ negotiation payoffs are regular and given by

$$r_i = \frac{1}{2}\left(\Pi(e(d)) - d_i + d_j\right)$$

where $e(d)$ is the best implementable effort vector given d and Π are the joint stage game payoffs. If$\rho > 0$ and ceteris paribus δ and $\bar{d}$ are sufficiently large, there won't be any incentive compatible relational contract in which any player makes herself unilateral vulnerable, since the high vulnerability $\bar{d}$ reduces her negotiation payoffs too strongly in future negotiations, which have sufficient weight for large δ. Even if both players plan to make themselves simultaneously unilaterally vulnerable, players have incentives to deviate. On the other hand, for sufficiently low implementation cost ε, it is always incentive compatible to propose mutual vulnerability and that is strictly preferred by both players compared to implementing no vulnerability at all due to the efficiency gains by being able to implement higher effort levels.

Proof of Proposition 5

In state x_1, negotiation payoff of player 1 is equal to $\beta_1 d$ if it holds that $c(1-\tilde{\delta}) \leq d\tilde{\delta}$. This player's payoff when state x_1 is reached then is at least $\frac{\rho\delta}{1-\tilde{\delta}}\beta_1 d$. It must be the case that weapons are acquired in a RNE if $\frac{\rho\delta}{1-\tilde{\delta}}\beta_1 d\phi > (1-\delta)b$ and $c(1-\tilde{\delta}) \leq d\tilde{\delta}$, which can only hold simultaneously for intermediate values of ρ.

Proof of Proposition 6

The players are either in the original state x^0 in which they can create a surplus within the relationship, or they have taken the outside option and receive π^{oo} forever after. Let $S(e) = \pi_1^{io} + \pi_2^{io} + e - k(e)$. In a truncated game with negotiation payoffs r, taking the outside option is a PPE if for some player i

$$\pi_i^{oo} > \frac{(1-\delta)\pi_i^{io} + \delta\rho r_i}{1-\delta(1-\rho)}. \tag{14}$$

Note that for $\delta \to 1$ (and $\epsilon \to 0$) this converges to $\pi_i^{00} \geq r_i$.

Assume that taking the outside option is not credible. In this case, players punish with the action $e = 0$, but punishment payoffs are larger than π_i^{oo} because of renegotiation. An effort level e can be played in state x^0 in a canonical relational contract if the following joint incentive constraint holds:

$$\tilde{\delta}(S(e) - S(0)) \geq (1-\tilde{\delta})k(e), \tag{15}$$

which means that $\bar{U} = S(e^0)$, where e^0 maximizes $S(e)$ over all e that satisfy (15). Negotiation payoffs are regular and equal to $r_i^{io} = \pi_i^{io} + \beta_i(S(e^0) - S(0))$. Taking the outside option is indeed not credible if for one i

$$(\pi_i^{oo} - \pi_i^{io})(1-\delta(1-\rho)) \leq \delta\rho\beta_i(S(e^0) - S(0)),$$

which never holds for β_i or ρ equal to zero. For $\delta \to 1$ it becomes $\pi_i^{oo} \leq r_i^{io}$. If this does not hold, then it must be the case that players take the outside option. In particular, if $\beta_i = 0$, then player i's payoff is equal to π_i^{00}.

Assume now that taking the outside option is credible, then an effort level e can be played in state x^0 in a canonical relational contract if

$$\delta(S(e) - (\pi_1^{oo} + \pi_2^{oo})) \geq k(e)(1-\delta), \tag{16}$$

which means that $\bar{U} = \max(S(e^{oo}), \Pi^{oo})$, where e^{oo} maximizes $S(e)$ over all e that satisfy (16). It then holds that $r_i = \pi_i^{00} + \beta_i(\bar{U} - \pi_1^{oo} - \pi_2^{oo})$. It is indeed credible to take the outside option if

$$(\pi_i^{oo} - \pi_i^{io})(1-\delta) > \delta\rho\beta_i(\bar{U} - \pi_1^{oo} - \pi_2^{oo}). \tag{17}$$

In particular, this holds for $\rho = 0$.

Proof of Proposition 7

In the truncated game with negotiation payoffs $r(x_0)$, a simple equilibrium can have $a^2 = a_\phi$ if

$$(1-\phi) \geq v_1(x_0|r) + v_2(x_0|r), \tag{18}$$

with the punishment payoff of the target being

$$v_2(x_0|r) \quad = \quad \frac{(1-\phi)(1-\delta+\delta\rho r_2)}{1-(1-\phi)\delta(1-\rho)} \tag{19}$$

and the punishment payoff of the blackmailer

$$v_1(x_0|r) = \frac{\delta\rho r_1}{1-\delta(1-\rho)}. \tag{20}$$

Hence the condition above is

$$(1-\rho)(1-\delta(1-\rho))(1-\phi) \quad \geq \quad \rho r_1. \tag{21}$$

When it is the target's turn to suggest a relational contract, she will unambiguously prefer a larger negotiation payoff. When it is the blackmailer's turn to select a relational contract, he prefers a larger negotiation payoff but only as long as punishment is credible. Regular negotiation payoffs constitute an upper bound on feasible negotiation payoffs:

$$r_1 \leq v_1 + \beta_1(1 - v_2 - v_1), \tag{22}$$

which here means

$$r_1 \leq \frac{\beta_1(1-\delta(1-\rho))\phi}{(1-\delta)(1-(1-\phi)\delta(1-\rho))+\beta_1\delta\rho\phi}. \tag{23}$$

The renegotiation payoff of an RNE has to be the maximum possible one. This means that if

$$(1-\delta)(1-\rho)(1-\phi) \geq \rho\beta_1\phi, \tag{24}$$

it holds that $r_1 = \frac{\beta_1(1-\delta(1-\rho))\phi}{(1-\delta)(1-(1-\phi)\delta(1-\rho))+\beta_1\delta\rho\phi}$. In that case, the RNE has regular negotiation payoffs, and the blackmailer's payoff is increasing in β_1 and ϕ. Intuitively, larger values of ϕ allow to extract larger payments in a subgame perfect equilibrium, and larger values of β_1 makes new negotiations more valuable.

However, if condition (24) does not hold, then it must be that $r_1 = \frac{(1-\rho)}{\rho}(1-\delta(1-\rho))(1-\phi)$. The blackmailer's negotiation payoff in this case is independent of his bargaining power β_1, since it is optimal to not exploit the negotiation opportunities too much anyway. Negotiation payoff is decreasing in δ, ρ and ϕ, since the punishment action is only credible if negotiation payoff is small. The optimal value of ϕ for the blackmailer is the one that makes condition (24) just bind.

Appendix B: Simple equilibria in truncated games

While there is no simple general recipe for finding canonical RNE, truncated games are also stochastic games, so that results from Goldlücke and Kranz (2017) apply. Some are stated here for the convenience of the reader.

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. Afterwards play can be either in the *equilibrium phase*, indexed by $k = e$, or in the *punishment phase* of some player i, indexed by $k = i$. A simple strategy profile specifies for each phase $k \in \mathcal{K} = \{e, 1, ..., n\}$ and state x an action profile $\boldsymbol{\alpha}^k(x) \in \mathcal{A}(x)$. We refer to $\boldsymbol{\alpha}^e$ as the equilibrium phase policy and to $\boldsymbol{\alpha}^i$ as the punishment policy for player i and call the vector of all policies $(\boldsymbol{\alpha}^k)_{k\in\mathcal{K}}$ an action 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. 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.

Consider now stochastic game with perfect monitoring. For a truncated game with expected negotiation payoffs r, let $U(x|\boldsymbol{\alpha}, r)$ denote the expected joint continuation payoffs (summing over payoffs for all players) in state x of a simple equilibrium with equilibrium phase policy $\boldsymbol{\alpha}$ and no money burning. These joint equilibrium payoffs can be easily computed by solving the following system of linear equations:

$$U(x|\boldsymbol{\alpha}, r) = (1-\delta)\Pi(x, \boldsymbol{\alpha}) + \delta E[(1-\rho)U(x'|\boldsymbol{\alpha}, r) + \rho R(x')|\boldsymbol{\alpha}, x], \tag{25}$$

where Π and R denote joint stage game payoffs and joint expected negotiation payoffs, respectively. To simplify notation, we abbreviate the joint payoffs in the equilibrium phase by

$$U^e(x) \equiv U(x|\boldsymbol{\alpha}^e, r).$$

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

$$v_i(x|\boldsymbol{\alpha}^i, r) = \max_{\hat{a}_i \in A_i(x)} \{(1-\delta)\left(\pi_i(\hat{a}_i, \boldsymbol{\alpha}^i_{-i}, x)\right) + \delta E[(1-\rho)v_i(x'|\boldsymbol{\alpha}^i, r) + \rho r_i(x')|\hat{a}_i, \boldsymbol{\alpha}^i_{-i}, x]\}. \tag{26}$$

It characterizes player i's best-reply payoffs in the truncated game if we assume that other players' action plan is fixed to $\boldsymbol{\alpha}^i_{-i}$ and no transfers are conducted.

Proposition 8. *Consider a stochastic game with perfect monitoring and let $(\boldsymbol{\alpha}^k)_k$ be an action plan in which in every state at least one player plays a pure strategy. A simple equilibrium with action plan $(\boldsymbol{\alpha}^k)_k$ exists for the truncated game with negotiation payoffs*

r if and only if for every state $x \in X$ and every phase $k \in \{e, 1, ..., n\}$

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

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:

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

The analysis of RNE often facilitates for the limit case $\delta \to 1$ with a fixed positive negotiation probability $\rho > 0$. Furthermore, in most games there won't be any money burning in the equilibrium phase or in negotiations and we have $U^e(x) = R(x)$ in all states x. The constraint (IC-SUM) then simplifies to

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