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

Susanne Goldlücke† and Sebastian Kranz‡

May 2022

Abstract

Game-theoretic analysis of relational contracts typically studies Pareto optimal equilibria. We illustrate how this equilibrium selection rules out very intuitive hold-up concerns in stochastic games with long-term decisions. The key problem is that Pareto optimal equilibria, even if satisfying renegotiation-proofness, do not reflect plausible concerns of how today's actions affect future bargaining positions within the relationship. We propose and characterize an alternative equilibrium selection based on the notion that continuation play is repeatedly negotiated in a relationship. We illustrate with several examples how the concept naturally combines relational contracting and hold-up concerns.

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 possibly affect the way in which 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

*We would like to thank Daniel Barron, Mehmet Ekmekci, 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

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] Given the common motivation and economists' immense interest in both fields, it is also important to have comprehensive frameworks for a unified analysis of relational contracting and hold-up problems. While advances have been made for repeated games[3], a tractable framework for dynamic stochastic games with infinite horizon is still missing. We seek to fill this gap by introducing our concept of *repeated negotiation equilibrium* (RNE).

Relational contracts are typically formulated as perfect public equilibria (PPE) of infinitely repeated games with payoffs on the Pareto frontier. In order to model relationships with long-term investments and corresponding hold-up problems, we study stochastic games, in which the stage game can change over time in response to players' actions. As we will illustrate with a simple example in Section 2, relational contracts can easily overcome many hold-up problems by making future trade and bargaining outcomes dependent on the conducted investments. This reveals the stark contrasts between the behavioral assumptions of the hold-up and the relational contracting literature. This example also introduces the core idea of RNE in a simple setting that abstracts from the intricacies of the infinite horizon context.

In Sections 3 and 4, RNE are formally introduced and characterized. The idea behind repeated negotiation equilibrium is that at the beginning of a period new negotiations are initiated with some exogenous probability. If new negotiations take place, bygones are bygones and only the payoff-relevant state matters for the outcome. 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. Both Pareto optimal PPE and MPE are solution concepts that are widely applied in the analysis of repeated and stochastic games, but they make very different assumptions about the history dependence of strategies. With a positive probability of new negotiations, the RNE concept naturally allows for intermediate cases.

In Section 5, we illustrate with a series of applications how the model leads to more plausible

[1] The hold-up problem has received a lot of attention since Grout's (1984) classic article, which shows how firms under-invest in capital because labor unions appropriate a share of the generated surplus in subsequent wage negotiations. Hold-up concerns play a crucial role for the organization of production (Klein et. al. (1978), Williamson (1985), Hart and Moore (1988)).

[2] Self-enforced 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. Other applications of relational contracting include cartels (e.g. Harrington and Skrzypacz (2011)), international trade agreements (e.g. Bagwell and Staiger (1990), Klimenko, Ramey and Watson (2008)) and environmental treaties (e.g. Barrett (2005), Harstad (2012, 2016), Gjertsen et al. (2021)). In all these applications, bargaining plays an important role.

[3] The connection between existing theories of ongoing negotiations in repeated games and our model is discussed below.

predictions than the corner cases. RNE are designed for studying the role of hold-up concerns and bargaining power when decisions have long-term consequences. One example are decisions to make oneself more vulnerable in a long-term relationship. While Pareto optimal PPE predict that players make themselves immediately strongly vulnerable to improve relational incentives, RNE yield a more complex picture in which players increase vulnerabilities in small steps and only as long as efficiency gains are sufficiently large. We also revisit a classical question from bargaining theory about the role of inside and outside option. While in Pareto-optimal PPE only outside-option payoffs are relevant, RNE stress the importance of inside options. Another example is the effect of asset ownership on investment decisions. We apply the RNE concept to a relational hold-up model, where the difference to the literature is that in an RNE, negotiations apply in the same way to the distribution of the surplus and to the punishment in case one party deviated. Despite enabling harsher punishments, joint ownership is dominated by single ownership as in the static model. Finally, we explore an arms race, where the goal of a stronger bargaining position leads to costly aquisition and potential usage of weapons in RNE, which does not happen in Pareto-optimal PPE or Markov perfect equilibria.

Since RNE do not exist in all stochastic games, we introduce an extension to weak RNE, which always exist, in Appendix A. Appendix B contains an algorithm to find RNE in games with perfect monitoring. All proofs are relegated to Appendix C.

Related literature

Whether hold-up problems arise in existing frameworks for renegotiation in infinitely repeated games depends critically on the nature of recurrent negotiation and the implications of disagreement. Renegotiation-proofness refinements for repeated games (e.g. Farrell and Maskin (1989), Bernheim and Ray (1989), Asheim (1991)) require that continuation equilibria are not Pareto dominated within a suitable set of equilibria. The idea behind these concepts is that the current agreement is the default which players return to under disagreement if one of the players blocks the negotiations. This Pareto criterion is criticized by Abreu, Pearce and Stacchetti (1993), who argue that renegotiation might be triggered more easily. The Pareto criterion is also relatively weak: In games with transfers, the use of fines often leads to Pareto efficient punishments, and therefore these concepts do not severely restrict the set of sustainable payoffs (Baliga and Evans (2000), Levin (2003), Goldlücke and Kranz (2013)).

While none of these early concepts capture the history independence that is needed to study hold-up, Miller and Watson (2013) propose a theory of negotiation to study the role of bargaining power in infinitely repeated games. Their concept of "contractual equilibrium" is built on the premise that continuation play under disagreement should not vary with the way in which bargaining breaks down. In a contractual equilibrium negotiations take place every period. Relational incentives can be supported by history-dependent disagreement points that involve no current period transfers. In contrast, RNE have history-independent disagreement points that cannot be chosen to support relational incentives, and transfers are possible also under disagreement.

The idea that many relations are characterized by ongoing negotiations has been explored in a few related articles. In the finite horizon context, Watson (2013) defines and classifies solution concepts according to the effect of negotiations and assumptions about disagreement. In the most flexible theory discussed by Watson (2013), the disagreement point can be a function of the players' shared history, as is assumed by Miller and Watson (2013). In the most stringent theory, the disagreement point does not depend on the history of play as in a repeated negotiation equilibrium. It becomes clear that the predictions of a bargaining theory crucially depend on the assumptions on disagreement and the degree of history-independence in the bargaining game. The contribution of RNE is to provide a general and tractable framework with a flexible specification of disagreement payoffs. Moreover, with the negotiation probability it introduces a parameter that measures the importance of history-independent bargaining power.[4]

While existing concepts are typically defined for repeated games, RNE applies to a more general class of dynamic games. However, Watson (2013) and Watson, Miller and Olsen (2020) distinguish between internally and externally enforced agreements, where the latter could alternatively be modeled as moving to another state of a stochastic game. Similarly, there is an earlier strand of the literature on renegotiation that studies dispute institutions that can also be viewed as states of the game (Ramey and Watson (2002), Klimenko et al. (2008)). These papers solve for a "recurrent agreement", which is characterized by Nash bargaining where under disagreement players expect to play a stage-game Nash equilibrium and resume negotiations in the following period. Klimenko et al. (2008) use the theory to explain how a dispute resolution mechanism can provide a benefit to the participating countries simply by allowing them to initiate a costly process following a deviation from cooperation. Such an institution could be studied with RNE in a model with two states, where moving from the initial state to the other (dispute settlement) state has to be an equilibrium.

Related is also a small literature that studies hold-up in repeated interactions. To address specific questions regarding for example the effect of asset ownership on investment, these models impose some bargaining structure that is tailored to the particular application. For example, Halac (2015) considers long-term investments that take place only in the first period and shows how the ensuing relational contract can mitigate the hold-up problem. Che and Sákovics (2004, 2021) and Pitchford and Snyder (2004) study hold-up in a dynamic setting and

[4]The negotiation probability can be interpreted as a measure of limited commitment. This is reminiscent of work on international agreements, where it is unclear how much commitment (legal or otherwise) there is to contractual agreements. For instance, Harstad (2012, 2016) introduces the length of a commitment period as an institutional parameter measuring the strength of the international contractual environment and investigates how it influences the hold-up problem.

show that making investment gradual can help to improve efficiency. 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.

2 Introducing repeated negotiation equilibria with a simple investment game

This section motivates repeated negotiation equilibria with a classic hold-up example. The game starts in an initial state x_0 in the first period where a buyer (player 1) and a seller (player 2) each choose investments a_i from a compact set $A_i(x_0)$, $i = 1, 2$. Investment costs for player i are given by $c_i(a_i) \geq 0$. Investments determine, possibly stochastically, the state x in period 2, which determines the total surplus achievable by trade $S(x)$. Payoffs are discounted between periods with a discount factor $\delta \in (0, 1)$. First best investments $a^* = (a_1^*, a_2^*)$ maximize the sum of expected payoffs given that ex post trade takes place whenever it is efficient,

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

In period 2 each player decides whether or not she wants to trade. If both players want to trade, each player gets a payoff of $\frac{1}{2}S(x)$. Otherwise each player gets an outside-option payoff of zero. From period 3 onward, the game is in an absorbing state x_{end} in which all players get zero payoff forever.[5]

Before actions are chosen, each period starts with a transfer stage in which each player can voluntarily transfer money to the other player (or burn money). Transfers are perfectly observed and strategies can condition actions on the observed transfers. Players are risk-neutral and net payments are simply added to players' payoffs in that period.

The voluntary transfers at the beginning of period 2 can be used to implement as a subgame perfect continuation payoff in state x every payoff u that satisfies $u_1 + u_2 \leq \max\{S(x), 0\}$ and grants each player at least the zero outside-option payoff.[6] That is because any deviation from a prescribed transfer can be punished by mutually choosing no-trade. The straight line segment in Figure 1 (left) illustrates the Pareto frontier of subgame perfect continuation equilibrium

[5] This game is in essence a two-period game, but we model it as an example of the general class of stochastic games studied in this article. In Section 5.3, we study an infinitely repeated version of this two-stage game.

[6] Being able to implement every split of the surplus as a subgame perfect continuation payoff is a fairly robust result. We get the same result if the formal contract specified a different split of the surplus or if trade decisions were modelled by a Nash demand game. Moreover, adding the possibility to come to an agreement in later periods if both player's reject trade in period 2 would not change the set of continuation payoffs in period 2. A famous exception is Rubinstein's (1982) alternative offer bargaining game which uniquely implements the Nash bargaining outcome. However, this unique equilibrium outcome is not robust with respect to plausible modifications of the bargaining game (e.g. Avery and Zemsky (1994)). Evans (2008) shows very generally that the multiplicity of equilibria of the bargaining game can be used to implement efficient investment.

payoffs in period 2 in a state x with positive surplus from trade.

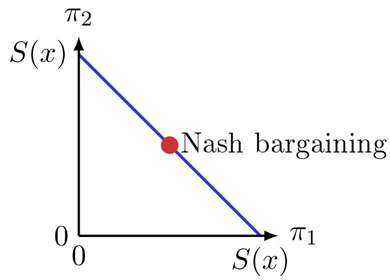

Figure 1: The blue line shows the Pareto-frontier of subgame perfect continuation payoffs in period 2. The thick line segment in the right figure illustrates the subset that can be implemented as expected continuation payoffs for a negotiation probability $\rho = 0.6$.

As a consequence, first best investments a^* are always conducted in a Pareto-optimal subgame perfect equilibrium (SPE). A player who unilaterally deviates from a^* will be punished with a continuation payoff of zero in all states in period 2. If no player unilaterally deviates, continuation equilibria split the surplus $S(x)$ such that on average each player gets at least her cost $c_i(a_i^*)$ reimbursed. Since the expected discounted joint surplus under first best investments is larger than total investment costs, such a split of trade surplus always exists. Key to this result is that in a Pareto optimal SPE continuation payoffs are flexibly picked depending on the conducted investments, as is typically the case in a relational contract.

In contrast, the hold-up literature commonly assumes that surplus from trade is always split according to the (symmetric) Nash bargaining solution, which in our example corresponds to an equal split of $S(x)$. As players do not receive the full return to their investment, underinvestment can arise. An important source of the hold-up problem is that writing a formal contract that conditions the split of the surplus on conducted investments is not possible because of a lack of external enforcement or verifiability.[7] Our example illustrates that in addition to the incompleteness of formal contracts, a limitation to the scope of relational contracting can also be essential for hold-up problems to arise.

Thus, a crucial difference between relational contracting and hold-up lies in whether or not continuation equilibria can flexibly depend on past actions. We do not attempt to answer the question which of the two ideas is the more appropriate concept. Experimental findings by Ellingsen and Johannesson (2004) and Ellingsen and Johannesson (2005), who compare behavior in a hold-up setting with either a Nash demand game or an ultimatum bargaining

[7] A particular focus of the literature is under which conditions simple contracts can reestablish efficient investments. 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 threat point in renegotiations (Aghion, Dewatripont and Rey (1994), Nöldeke and Schmidt (1995), Edlin and Reichelstein (1996)). If the seller's investments can influence the buyer's valuation and vice versa, Che and Hausch (1999) show that the hold-up problem cannot be resolved by any contract on output that is renegotiated according to the Nash bargaining solution.

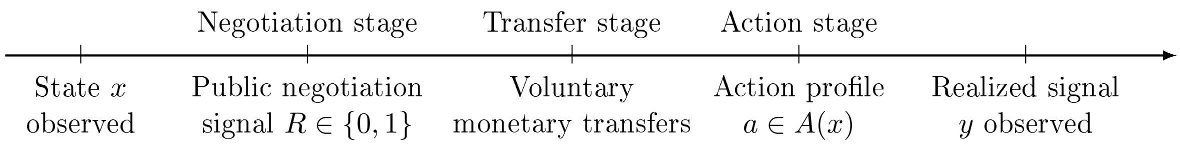

Figure 2: Timeline of a period

game, support the view that intermediate cases are plausible.[8] The model we propose provides a framework that unites both ideas by introducing a continuum of intermediate cases.

In an RNE we assume that at the beginning of each period, with some exogenous negotiation probability $\rho \in [0, 1]$ new negotiations take place and the players receive a continuation payoff that corresponds to the Nash bargaining solution over the set of continuation equilibrium payoffs (holding future negotiations fixed). Assuming no-trade as disagreement point and equal bargaining weights, new negotiations in period 2 would mean that the surplus $S(x)$ is split equally. Consider the case that a player has deviated from required investments and is supposed to be punished by zero continuation payoffs. Given the possibility of negotiation in period 2, the deviating player is then still able to guarantee herself an expected continuation payoff of $0.5\rho S(x)$. Hence, the span of Pareto optimal expected continuation payoffs that can be implemented in state x is a fraction $1 - \rho$ of the span of the Pareto optimal continuation payoffs. Figure 1 (right) shows the range of implementable expected payoffs for $\rho = 0.6$. Above a critical negotiation probability, it is no longer possible to implement first best investments.

3 Model

We consider n-player stochastic games with infinitely many periods. 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 three stages: a negotiation stage, a transfer stage and an action stage. There is no discounting between stages. Figure 2 illustrates the timeline of a period.

At the beginning of each period t, new negotiations take place with an exogenously given probability $\rho \in [0, 1]$. This is modeled by an exogenous public signal $R_t \in \{0, 1\}$ that the players observe in the negotiation phase. A positive negotiation signal $R_t = 1$ indicates new negotiations while $R_t = 0$ indicates no new negotiations. We assume that the initial period $t = 0$ always starts with negotiations, i.e., $R_0 = 1$. Note that the signal R_t will only be used in the definition of an RNE to trigger continuation equilibria that are interpreted as the outcome of new negotiations, but otherwise has no direct payoff consequence.

In the transfer stage, every player simultaneously chooses a non-negative vector of transfers

[8] See also Ellingsen and Robles (2002) and Tröger (2002) for evolutionary arguments why bargaining outcomes can depend on prior investment.

to all other players. 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 an action from a finite or compact action set $A_i(x)$. The set of (pure) action profiles in state x is denoted by $A(x) = A_1(x) \times ... \times A_n(x)$.

We allow for imperfect public monitoring. After actions have been conducted, a signal y from a finite signal space Y and then the next period's state are commonly observed. Both are jointly drawn from a probability distribution that only depends on the current state x and action profile a. Player i's stage game payoff is denoted by $\hat{\pi}_i(x, a_i, y)$ and depends on the initial state x, player i's action a_i, and the signal y. In a game with perfect monitoring, the signal y is equal to the played action profile a. We denote by $\pi_i(x, a)$ player i's expected stage game payoff in state x if action profile a is played. 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. The expected payoff of player i in a period in state x, in which she makes a net transfer of p_i and action profile a is played, is then given by $\pi_i(x, a) - p_i$. When referring to payoffs of the stochastic game, we mean expected average discounted payoffs, i.e., the expected sum of discounted payoffs multiplied by $(1 - \delta)$.

We either restrict attention to pure strategies or, for finite action spaces, also consider behavioral strategies in which players can mix over actions. If behavioral strategies 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 $\pi_i(x, \alpha)$ player i's expected stage game payoff taking expectations over mixing probabilities and signal realizations.

A public history h of the stochastic game is a sequence that specifies up to a particular point of play all states, negotiation signals, performed transfers and public signals that have so far occurred. A public strategy σ_i of player i maps every public history that ends before a payment stage into a vector of monetary transfers to other players, and every public history ending before the action stage into a (possibly mixed) action $\alpha_i \in \mathcal{A}_i(x)$. Let Σ denote the set of public strategy profiles. A public perfect equilibrium (PPE) is a profile of public strategies that constitute mutual best replies after every history.

Given a public history h, we denote by $h^C(h)$ the ending subsequence of h that starts with the latest period in which negotiations took place in h. This means that only the first negotiation signal in the history $h^C(h)$ is equal to 1. We define by

$$\Sigma^* = \left\{\sigma \in \Sigma | \ \sigma(h) = \sigma(h') \text{ if } h^C(h) = h^C(h')\right\} \tag{1}$$

the subset of all public strategy profiles in which the chosen actions and transfers only depend on the history since the last negotiations. In particular, directly after new negotiations take place, i.e. if $h^C(h) = (x, 1)$, continuation play of a strategy profile in Σ^* only depends on the current state x. For any $\sigma \in \Sigma^*$, we denote by $r^\sigma(x)$ the well-defined profile of continuation payoffs directly after new negotiations take place in state x. We call r^σ the negotiation payoffs of σ.

A repeated negotiation equilibrium will be defined as a strategy profile $\sigma \in \Sigma^*$ that is a PPE and whose negotiation payoffs correspond for each state to a generalized Nash bargaining solution over a suitably defined bargaining set. Key for the specification of the bargaining set is to characterize the set of possible continuation payoffs taking future negotiation payoffs as given. For this characterization, we consider *truncated games*, which end with some fixed payoff vector when new negotiations are triggered.

A truncated game $\Gamma(x, r)$ is defined by a state x and an arbitrary vector of (negotiation) payoffs r that specifies for every state $x' \in X$ a payoff profile $r(x') \in \mathbb{R}^n$. The truncated game $\Gamma(x, r)$ is equal to the original game with the following modifications: It starts in state x and whenever play would transit to a state x' then with probability ρ an absorbing state is reached in which players receive the terminal payoffs $r(x')$ in the current and every future period and no more actions and transfers are possible. Otherwise, action spaces, state transitions and payoffs of the truncated game are the same as in the original game. The truncated game always starts with at least one period of regular play before an absorbing state can be reached. For any strategy profile σ of the truncated game $\Gamma(x, r)$, we denote player i's payoff by $u_i(\sigma, x, r)$.

We can naturally map each strategy profile $\sigma \in \Sigma^*$ of the original game to a tuple $(\sigma^x)_{x \in X}$ of public strategy profiles of the truncated games $\Gamma(x, r^\sigma)$ where σ^x follows the transfers and actions of σ starting from new negotiations in state x until again new negotiations take place.

Lemma 1. *A strategy profile $\sigma \in \Sigma^*$ with negotiation payoffs r^σ is a PPE in the original game if and only if for its representation as strategy profiles $(\sigma^x)_{x \in X}$ in the truncated games $\Gamma(x, r^\sigma)$, each σ^x is a PPE of $\Gamma(x, r^\sigma)$ with payoff profile $u(\sigma^x, x, r^\sigma) = r^\sigma(x)$.*

As in every stochastic game with transfers, the PPE payoff set of a truncated game $\Gamma(x, r)$ is a simplex, which we denote by

$$\mathcal{U}(x, r) = \{u \in \mathbb{R}^n | \sum_{i=1}^{n} u_i \leq \bar{U}(x, r) \text{ and } u_i \geq \bar{v}_i(x, r) \text{ for all } i = 1, ..., n\}, \qquad (2)$$

where $\bar{U}(x, r)$ denotes the maximum joint PPE payoff and $\bar{v}_i(x, r)$ the minimum PPE payoff for player i.

To model negotiation outcomes that follow a positive negotiation signal as a generalized Nash bargaining solution, additional ingredients are needed: bargaining weights and disagreement points. Bargaining weights are given exogenously by $\beta_1, ..., \beta_n \in [0, 1]$ with $\sum_{i=1}^{n} \beta_i = 1$.

Disagreement points are specified by a function $d(x, r)$ that assigns to every state x a payoff profile for any given vector of negotiation payoffs r. The crucial aspect is that the disagreement point only depends on future negotiation payoffs and the current state x, but not on any other aspect of the history. Our model allows for different specifications of the disagreement points, which shall satisfy

Assumption 1. *For any truncated game $\Gamma(x, r)$ it must hold that $d(x, r) \in \mathcal{U}(x, r)$. Moreover, if for truncated games $\Gamma(x, r)$ and $\Gamma(x, r')$ there is a vector $b \in \mathbb{R}^n$ such that $u(\sigma, x, r) = u(\sigma, x, r') + b$ for every strategy profile σ of the truncated game that starts in state x, then $d(x, r) = d(x, r') + b$.*

Hence, disagreement payoffs are required to be PPE payoffs in the truncated game. Moreover, if negotiation payoffs in a truncated game are changed in such a way as to only shift the payoffs of all strategies by some constant, then the disagreement points should shift in the same way.

Definition 1. A PPE $\sigma \in \Sigma^*$ is a repeated negotiation equilibrium (RNE) for bargaining weights $\beta_1, ..., \beta_n$ and disagreement point function $d(x, r)$ if its negotiation payoffs r^σ satisfy

$$r_i^\sigma(x) = d_i(x, r^\sigma) + \beta_i(\bar{U}(x, r^\sigma) - \sum_{j=1}^{n} d_j(x, r^\sigma)) \text{ for all } x \in X \text{ and } i = 1, .., n. \tag{3}$$

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, with the additional assumption that at the beginning of the game, players receive their Nash bargaining payoff. If the negotiation probability is one and the game has a unique MPE payoff, then the RNE is an MPE.

Disagreement points

While our concept allows for different specifications of disagreement points, we will assume in our applications that the disagreement point $d(x, r)$ corresponds to the lowest PPE payoffs $\bar{v}(x, r)$ of the truncated game $\Gamma(x, r)$. This specification is particularly tractable since the computation of the PPE payoff set $\mathcal{U}(x, r)$ of the truncated game anyway requires to compute the corresponding lowest PPE payoffs. It is also intuitive in so far that in optimal equilibria of the truncated games, observable deviations will also be punished with minimal PPE payoffs, such that players threaten with similar punishment actions after deviations and disagreement.

The resulting negotiation payoffs are also consistent with a random dictator specification of disagreement proposed in the literature (see Watson, 2013). Translated to our framework, under disagreement each player i is a random dictator with probability β_i and chooses continuation play that grants all other players $j \neq i$ their minimum continuation payoff $\bar{v}_j(x, r)$

and him the surplus $\bar{U}(x,r) - \sum_{j \neq i} \bar{v}_j(x,r)$. Then $d(x,r)$ denotes the expected disagreement payoff. The resulting negotiation payoffs specified by (3) are the same as under our assumption $d(x,r) = \bar{v}(x,r)$.

Another idea (e.g. Miller and Watson, 2013) is that under disagreement no transfers are conducted. One could translate this assumption to our framework using a random dictator formulation in which only continuation equilibria of the truncated game without transfers can be chosen under disagreement. However, that assumption would make characterization of RNE considerably more complex and it can become quickly intractable in stochastic games with several states.[9]

Another tractable and in some applications very natural assumption on disagreement points is that temporarily a stage game Nash equilibrium is played (Ramey and Watson, 2002). In our framework, this is most closely matched by assuming that disagreement payoffs correspond to MPE payoffs of the truncated game. In most of our applications, worst punishment payoffs $\bar{v}(x,r)$ indeed also constitute MPE payoffs. Yet, the arms race application in Section 5.4 illustrates different predictions of the two concepts.

Endogenous negotiations

RNE can be interpreted as outcomes of endogenous negotiations in the following way. In the negotiation stage each player can always attempt to force new negotiations of the relational contract. $R_t = 1$ then indicates that there was a successful negotiation attempt while $R_t = 0$ indicates that there was no negotiation attempt or no attempt was successful. A negotiation attempt shall only be successful with an exogenously given probability ρ, and if it is not successful, it is not observed. This assumption rules out that players can be punished for unsuccessful negotiation attempts.[10] Neither are players punished for successful negotiation attempts, which yield the negotiation payoffs. Since negotiation payoffs are Pareto optimal continuation payoffs (taking future negotiation outcomes as given), either all players are indifferent between negotiating or not, or there is always a player who strictly prefers to attempt new negotiations. Assuming that also indifferent players attempt new negotiations, we get the same outcome as in the RNE: negotiations take place with probability ρ each period.

Of course, ex-ante players would typically benefit from being able to commit to never initiate such renegotiation attempts. With this alternative model, $1 - \rho$ can be interpreted as an exogenous measure of the institutional strength to ignore such renegotiation attempts.

[9] A more important difference to Miller and Watson (2013) is that in RNE disagreement payoffs must only depend on the current state and not on any other aspect of the history. The application in Section 5.1 sheds additional light on this point.

[10] See for example Safronov and Strulovici (2018) for a repeated game model in which players might be deterred from proposing new negotiations. In a setting with transfers, punishing a player for a proposal would always be feasible if the consequences of disagreement could depend on the identity of the proposer.

4 Characterization of RNE

4.1 Simple RNE

An important step when characterizing RNE is to find PPE payoff sets of truncated games. Goldlücke and Kranz (2018) show that every PPE payoff of a stochastic game with transfers can be implemented with a class of simple equilibria. For convenience, we give a brief summary here.

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 $\alpha^k(x) \in \mathcal{A}(x)$. We refer to α^e as the equilibrium phase policy and to α^i as the punishment policy for player i. 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. A simple PPE has a simple strategy profile. There is always an optimal policy $\{\alpha^k\}_k$ such that every PPE payoff can be implemented with a simple PPE with that optimal policy by varying first period's upfront transfers (possibly burning money).

In a simple strategy profile with negotiations, transfers can in addition condition on the negotiation signal: Whenever new negotiations take place in state x, players perform transfers defined by a vector of net transfers $p^0(x)$ that only depends on the state x. A simple RNE shall be such a simple strategy profile with negotiations that is a RNE. Moreover, a simple RNE's policy $\{\alpha^k\}_k$ shall be an optimal policy in all truncated games. Since negotiations affect the path of play only by modifying the subsequent upfront payments, a simple RNE has a particularly tractable structure. Conveniently, we can restrict attention to simple RNE.

Proposition 1. *For every RNE there exists a simple RNE with the same negotiation payoffs.*

4.2 Computing pure strategy, simple RNE in games with perfect monitoring

We now present a brute-force method to compute all simple pure strategy RNE in a game with perfect monitoring given the assumption that disagreement payoffs are equal to worst punishment payoffs. One iterates through all possible policies $\{\alpha^k\}_{k=e,1,..,n}$ as candidate policy for a simple RNE. For any state x, the highest joint equilibrium payoff that can be implemented with α^e is computed as

$$U(x|\alpha^e) = (1-\delta)\Pi(x,\alpha^e) + \delta E[U(x'|\alpha^e)|x,\alpha^e], \tag{4}$$

where Π is the joint stage game payoff and the expectation is taken over the transition probabilities to the next period's state x'. For given negotiation payoff vector r the lowest punishment payoffs that can be imposed on player i given a punishment policy α^i are characterized by the solution of the following Bellman equation:

$$v_i(x|\alpha^i,r) = \max_{\hat{a}_i \in A_i(x)} \{(1-\delta)\pi_i(x,\hat{a}_i,\alpha^i_{-i}) + \delta E[(1-\rho)v_i(x'|\alpha^i,r) + \rho r_i(x')|x,\hat{a}_i,\alpha^i_{-i}]\}. \tag{5}$$

Given the assumption that under disagreement optimal punishment payoffs are implemented, we can use (3) to substitute the negotiation payoffs r. Punishment payoffs thus satisfy:

$$\begin{aligned} v_i(x|\alpha) = \max_{\hat{a}_i \in A_i(x)} &\{(1-\delta)\pi_i(x,\hat{a}_i,\alpha^i_{-i}) + \\ &\delta E[v_i(x'|\alpha) + \rho(\beta_i(U(x|\alpha^e) - \sum_{j=1}^{n} v_j(x'|\alpha)))|x,\hat{a}_i,\alpha^i_{-i}]\}. \end{aligned} \tag{6}$$

This system of linear Bellman equations can be jointly solved for all players. Resulting negotiation payoffs are simply computed as

$$r_i(x) \quad = \quad v_i(x|\alpha) + \beta_i(U(x|\alpha^e) - \sum_{j=1}^{n} v_j(x,\alpha)) \text{ for all } i = 1,..,n. \tag{7}$$

It then follows from Goldlücke and Kranz (2018, Theorem 2) that a simple SPE with policies $(\alpha^k)_k$ exists in the corresponding truncated game if and only if for every state $x \in X$ and every phase $k \in \{e,1,...,n\}$ the following joint incentive constraint is satisfied:

$$\begin{aligned} &(1-\delta)\Pi(x,\alpha^k) + \delta E\left[U(x'|\alpha^e)|x,\alpha^k\right] \geq \\ &\sum_{i=1}^{n} \max_{\hat{a}_i \in A_i(x)} \left\{(1-\delta)\pi_i(x,\hat{a}_i,\alpha^k_{-i}) + \delta E\left[(1-\rho)v_i(x'|\alpha^i,r) + \rho r_i(x')|x,\hat{a}_i,\alpha^i_{-i}\right]\right\}. \end{aligned} \tag{8}$$

If (8) is violated, no simple RNE with policy $\{\alpha^k\}_{k=e,1,..,n}$ exists. Otherwise, it is neccessary and sufficient that the candidate policy implements the highest joint SPE payoffs and worst SPE punishment payoffs in the truncated game with negotiation payoffs r. To check this condition, one could run a similar inner loop through all possible policies. As in the outer loop, one uses conditions (4) and (5) to compute equilibrium path and punishment payoffs and then checks the incentive constraints (8). The only difference is that in this inner loop negotiation payoffs r are fixed. Alternatively, one could use the faster algorithm from Goldlücke and Kranz (2018) to compute the set of SPE payoffs in the truncated game.

We will illustrate in our applications that the procedure above can often be simplified, e.g. because one can easily narrow the set of candidate policies or symmetry considerations can pin down negotiation payoffs. In Section 4.4 we introduce the related concept of T-RNE, which can be far more efficiently computed numerically than RNE in large games.

4.3 Non-existence: Gangsta's Paradise game

The following example has two purposes: it showcases how to compute RNE using the conditions above and it illustrates that RNE can fail to exist. The *Gangsta's Paradise* game has two states: a prisoners' dilemma state x_0 and paradise. In x_0 players play a prisoners' dilemma game with the following payoff matrix:

	C	D
C	2,2	-10,3
D	$3,-10$	0,0

Players move to paradise if and only if both players cooperate. In paradise, players get a fixed payoff of $(2,2)$ without performing any actions. Paradise is an absorbing state in which players stay forever. Players shall have equal bargaining weight and mutually defect under disagreement in x_0. We will show that there is a range of discount factors and negotiation probabilities in which this Gangsta's Paradise game has no RNE.

First consider a candidate for a simple RNE in which players mutually cooperate on the equilibrium path so that for each player i negotiation payoffs satisfy $r_i(x_0) = 2$. The Bellman equation for the punishment payoffs (5) becomes

$$v_i(x_0) = \delta[(1-\rho)v_i(x_0) + \rho r_i(x_0)] \tag{9}$$

which yields

$$v_i(x_0) = \frac{2\delta\rho}{1-\delta(1-\rho)}. \tag{10}$$

If we want to implement mutual cooperation, the incentive constraint (8) becomes in this symmetric situation

$$2 \geq (1-\delta)\cdot 3 + \delta[(1-\rho)v_i(x_0) + \rho r_i(x_0)]. \tag{11}$$

and simplifies to

$$2 \geq (1-\delta)\cdot 3 + v_i(x_0). \tag{12}$$

If and only if this incentive constraint is satisfied, an RNE with mutual cooperation exists. Assume now (12) is violated. Consider a candidate for a simple RNE in which players mutually defect in state x_0, yielding $r_i(x_0) = v_i(x_0) = 0$. This candidate can only be an RNE if the more efficient action profile (C, C) cannot be implemented in the corresponding truncated game.[11] For an optimal simple SPE with mutual cooperation in the truncated game, joint equilibrium payoffs according to (4) are simply $U(x_0) = 4$, since after mutual cooperation paradise is always reached. The incentive constraint (8) to implement mutual cooperation in this truncated game becomes

$$2 \geq (1 - \delta) \cdot 3. \tag{13}$$

Note that (13) is satisfied for a larger set of δ and ρ combinations than (12). That is because lower negotiation payoffs in state x_0 reduce incentives to deviate from mutual cooperation since staying in state x_0 becomes less attractive. Assume e.g. $\delta = 0.5$ for which condition (13) is strictly satisfied. But for $\rho > \frac{1}{3}$ condition (12) is strictly violated and no RNE exists. The problem is that for those parameters mutual cooperation can be implemented today if and only if players expect that mutual cooperation is not implemented in future negotiations. Note that the non-existence result is not due to an unusual selection of disagreement payoffs, as mutual defection is an MPE of the original game and also yields the minimum payoff. The absorbing state is not the reason behind the non-existence result, either. We could alternatively assume that with some exogenous probability $\varepsilon > 0$ play transits from paradise back to x_0. If ε is not too large, the existence problem prevails, although for a smaller parameter range. In the remainder of this section we identify classes of games for which RNE always exist and have a unique payoff. Moreover, in Appendix A we introduce and illustrate weak RNE, which always exist.

4.4 Repeated games

A repeated game corresponds to the special case of a stochastic game with a single state. The set of PPE payoffs of a repeated game with transfers for discount factor δ is given by

$$\mathcal{U}(\delta) = \{u \in \mathbb{R}^n | \sum u_i \leq \bar{U}(\delta) \text{ and } u_i \geq \bar{v}_i(\delta) \text{ for all } i\}, \tag{14}$$

where $\bar{U}(\delta)$ is the highest joint PPE payoff and $\bar{v}_i(\delta)$ is the lowest PPE payoff for player i. We will show that in a repeated game, new negotiations have a similar effect as a restart of the relationship, such that a positive negotiation probability essentially reduces the effective discount factor to an adjusted discount factor $\tilde{\delta} = (1 - \rho)\delta$. Since a repeated game has only a single state, we suppress the dependence on the state in the notation, such that $d(0)$ denotes the profile of disagreement payoffs for the truncated game with zero negotiation payoffs.

[11]While also mixed action profiles may be more efficient than (D, D) given a sufficient a high probability to cooperate, they can be neglected because they are even harder to implement than (C, C).

Proposition 2. *In a repeated game, an RNE always exists and RNE payoffs are uniquely given by*

$$r_i = \tilde{d}_i + \beta_i(\bar{U}(\tilde{\delta}) - \sum_{j=1}^{n} \tilde{d}_j), \tag{15}$$

where $\tilde{\delta} = \delta(1-\rho)$ is an adjusted discount factor and $\tilde{d} = \frac{(1-\tilde{\delta})}{(1-\delta)} d(0)$ is an adjusted disagreement payoff. If one assumes that under disagreement players continue with a Nash equilibrium α^* of the stage game, then $\tilde{d}_i = \pi_i(\alpha^*)$. If ones assumes that players continue with worst punishments under disagreement, then $\tilde{d}_i = \bar{v}_i(\tilde{\delta})$.

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. Our result that in repeated games critical negotiation probabilities are basically equivalent to critical discount factors is reassuring since it means that the usual analysis of institutions and relational contracting is robust to the introduction of repeated negotiations. In contrast, in stochastic games first-best strategies may depend on the discount factor, so that it has anyway little appeal to study the minimal discount factors for which first-best strategies can be implemented.[12] First-best strategies are not affected by the negotiation probability, however, which makes the critical negotiation probability a more suitable measure to study comparative statics of relational contracting in stochastic games.

4.5 Strongly directional games and T-RNE

The RNE existence and uniqueness result for repeated games can be extended to an important subclass of stochastic games in which play will eventually reach an absorbing state and the stochastic game becomes a repeated game.

Definition 2. A stochastic game is *strongly directional* if each non-absorbing state can be visited at most once.[13]

Possible state transitions in a strongly directional game can be described by a directed acyclic graph, in which nodes represent the states and edges represent possible transitions from one state to another.

Proposition 3. *Any strongly directional game has an RNE and RNE payoffs are unique.*

[12] Consider the example in Section 2, in which players can perform costly investments that have positive externalities in the future. The critical discount factor to implement the first best is then trivially $\delta^* = 0$. That is because for $\delta = 0$, the first best solution, which is also the outcome in every SPE, is to never conduct any costly investments.

[13] We have adapted the terminology from Iskhakov et. al. (2015). They define *directional games* for which they develop an algorithm to quickly compute the set of all MPE. In a directional game, one can remain in a non-absorbing state for more than one period, which is not possible in a strongly directional game. Both definitions exclude cycles between multiple non-absorbing states.

The assumption that every state is only entered once avoids the recursive structure of negotiation payoffs and makes it possible to find an RNE via backward induction. If one restricts attention to pure strategies and games of perfect monitoring, there is an algorithm that can numerically compute RNE in a very quick fashion (see Appendix B).[14]

T-RNE

The following construction is very useful to effectively numerically study repeated negotiations in stochastic games that are not strongly directional. Take any stochastic game and fix some large period T, e.g. $T = 10000$. Consider the following derived strongly directional game. In period T the strongly directional game reaches an absorbing state (x, T) with $x \in X$ that has a truncated PPE payoff set equal to the PPE payoff set of the original stochastic game starting in state x.[15] A non-absorbing state consists of a tuple (t, x) of the current period $t < T$ and a state x of the original game. The component x transits as in the original game and t is incremented by 1 each period.
The interpretation of this strongly directional game is as following. The original stochastic game is played but new negotiations can occur only up to period T. Afterward no new negotiations take place, but the state can still change as in the original stochastic game. We call the RNE of this strongly directional game a *T-RNE* of the original game. Applying Proposition 3 we find

Corollary 1. *Consider a stochastic game, fix some period $T \geq 1$ and assume that new negotiations can occur only up to period T. A corresponding T-RNE always exists and we have unique T-RNE payoffs.*

5 Applications

5.1 Vulnerability and starting small in a principal-agent relationship

Consider a principal-agent relationship where in each period the agent (player 2) picks an effort level $e \in \{0, 0.01, ..., 1\}$ that benefits the principal (player 1) but is costly to the agent. Part of the state definition is for each player $i = 1, 2$ a value $x_i \in \{0, 0.1, ..., 0.5\}$ that shall measure how dependent player i is on the other player. One reason why a stronger dependency could be beneficial in real world relationships is that it may entail direct cost savings. Yet, a stronger dependency typically also makes a party more vulnerable. To highlight the strategic aspects of a stronger dependency, we abstract from any direct technological benefits. This

[14] We have implemented this and other algorithms in an R package that contains several examples, see `https://skranz.github.io/RelationalContracts`.

[15] Since the PPE payoff set for the original stochastic game is just a simplex, it is clear that one can always construct a repeated game for the absorbing state that has the same PPE payoff set. The exact structure of that repeated game is irrelevant, however.

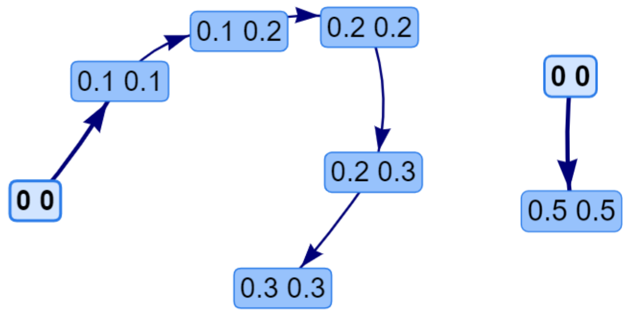

Figure 3: Transitions of vulnerabilities starting from $x_1 = x_2 = 0$ for $\tilde{\delta} = 0.15$ with $\rho = 0.5$ (left) and $\rho = 0$ (right).

means x_i shall only measure how vulnerable player i is towards the other player. Each period the principal can choose harm $h_1 \in \{0, x_2\}$ against the agent and the agent can choose harm $h_2 \in \{0, x_1\}$ against the principal. Resulting stage game payoffs are

$$\pi_1(e, h) = e - h_2,$$
$$\pi_2(e, h) = -\frac{1}{2}e^2 - h_1.$$

To guarantee existence and a unique RNE payoff, we assume that it is only possible to change vulnerabilities until period $T = 100$. Afterwards the game reaches absorbing states in which players face repeated games: vulnerabilities remain fixed at their current level while effort and harm levels can still be freely chosen every period. We thus have a strongly directional game in which a state in period $t \leq T$ is described by a tuple (x_1, x_2, t).

In a repeated game with fixed vulnerabilities x_1 and x_2 the implementable effort level only depends on the adjusted discount factor $\tilde{\delta} = (1 - \rho)\delta$, which means a higher negotiation probability is equivalent to a lower discount factor. Yet, in the stochastic game repeated negotiations can strongly affect the dynamics of endogenous vulnerabilities. Figure 3 shows the development of vulnerabilities and effort levels on the equilibrium path starting with $x_1 = x_2 = 0$. We fix the adjusted discount factor at $\tilde{\delta} = 0.15$ and compare a positive negotiation probability of $\rho = 0.5$ (left) with the case of no repeated negotiation $\rho = 0$ (right), which constitutes a Pareto optimal SPE.

We see that with repeated negotiations both players increase their vulnerabilities in small steps. Increasing vulnerability is beneficial since higher effort levels can be implemented with harsher punishment opportunities. Yet, it is not incentive compatible that both players immediately increase their vulnerability to the maximum level in one large step. The problem is that by deviating and keeping low vulnerability a player can keep a strong bargaining position that can be used to extort payments from the vulnerable player in future negotiations. In other words, a unilaterally vulnerable player faces a hold-up problem. Smaller steps can be incentivized since one can reward such a step with positive transfers and punish deviations if there are no new negotiations in the next period. Players increase their mutual vulnerability only up to intermediate levels $x_1 = x_2 = 0.3$. Even though higher vulnerabilities can incentivize even higher effort, there is the short run cost that one can incentivize less effort in a period where a player shall be incentivized to take an additional vulnerability step. Due to the convex cost function the gains from increasing vulnerabilities beyond 0.3 are just too small compared to those short run costs.

The diagram on the right shows that without repeated negotiations, in a Pareto optimal SPE, both players immediately make themselves maximally vulnerable. This is efficient as the highest implementable effort levels can be implemented as quickly as possible. It is also incentive compatible, because without repeated negotiations a player faces no hold-up problem when being more vulnerable: in a Pareto optimal SPE it is assumed that harm is only inflicted as a punishment while players can coordinate to ignore all attempts to exploit the vulnerability on the equilibrium path. Indeed, one can generally show that even if just a single player could make herself very vulnerable, she would immediately do so in a Pareto optimal SPE.

Vulnerabilities can stem from various sources, depending on the application. They can for example be the consequence of formal contracts, which –if enforced- would yield an outcome that one or both parties would like to avoid. In buyer-supplier relationships, 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.[16]

This example shows how RNE incorporate hold-up concerns in relational contracting. In contrast, the renegotiation-proofness concept of strong optimality (Levin, 2003) requires that all continuation payoffs lie on the Pareto frontier of SPE payoffs. Strong optimality therefore forces the Pareto optimal SPE outcome that both players make themselves immediately vulnerable in one large step. A direct comparison to contractual equilibria (Miller and Watson, 2013) is difficult as they are defined for repeated games only. However, a crucial aspect of contractual equilibria is that disagreement actions depend on the history, which means that players can coordinate to not exploit vulnerabilities under disagreement. Therefore, a Pareto criterion would suggest that also in a contractual equilibrium players would immediately choose the highest vulnerability level.[17]

Thus, a gradual increase in vulnerability is one testable implication of RNE not implied by

[16]Another reason to start small in relationships is that parties need time to learn each other's types in a framework of incomplete information, see Watson (1999) and Watson (2002).

[17]Note that although vulnerability in general helps to implement higher effort levels in a contractual equilibrium, it cannot substitute for complete lack of bargaining power: If the agent has zero bargaining power, the attainable surplus in a contractual equilibrium is zero independent of the vulnerability.

other concepts of renegotiation in repeated relationships. 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 prisoner's dilemma, then a gradual mechanism with incremental and conditional increases in vulnerability is more effective than simple simultaneous choices of the vulnerability level.

5.2 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 well-known outside option principle states that outside options should only influence bargaining outcomes if they are binding, while otherwise only inside options are relevant (see Binmore, Rubinstein and Wolinsky (1986) and 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)) or wage bargaining (Hall and Milgrom (2008)). 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. We assume again that disagreement payoffs are equal to the worst punishment payoffs, i.e. $d_i(x,r) = \bar{v}_i(x,r)$. Within the relationship, the principal's and agent's stage game payoffs shall be given by

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

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 end 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 $\pi_i^{io} < \pi_i^{oo}$, which means that both players prefer a break-up compared to staying in an unproductive relationship. To rule out that a player is indifferent between quitting or not if the other player wants to quit, we assume that if just one player wants to leave, there is a very small probability $\varepsilon > 0$ that the break-up is not successful and players remain in the relationship next period.

Proposition 4. *If the negotiation probability is zero, RNE payoffs are only determined by the outside options and uniquely given by*

$$r_i^{oo} = \pi_i^{oo} + \beta_i(\bar{U} - \pi_1^{oo} - \pi_2^{oo}), \tag{16}$$

where $\bar{U}$ denotes the joint payoff of the RNE. If instead $\rho > 0$, then in the limit $\delta \to 1$, RNE payoffs satisfy the outside option principle: Unless the outside option is binding for some player, they are solely determined by the inside options and given by

$$r_i^{io} = \pi_i^{io} + \beta_i(e^{io} - k(e^{io})) \text{ for } i = 1, 2, \tag{17}$$

where e^{io} is the optimal effort level in the repeated game without the outside options. If the outside option is binding for player i (meaning $\pi_i^{oo} \geq r_i^{io}$), then player i gets a negotiation payoff of π_i^{oo}.

First consider the case that the probability of future negotiations is zero, such that negotiation payoffs are described by (16). The disagreement payoffs are given by the lowest SPE payoff π^{oo}, implemented by the credible threat to choose the outside option under disagreement. With a larger negotiation probability choosing the outside option can cease to be incentive compatible. While a refusal to take the outside option can be punished until new negotiations take place, the prospect of sufficiently high negotiation payoffs can make it just too attractive to stay in the relationship.

This observation yields a simple intuition for the limit case $\delta \to 1$ with a fixed $\rho > 0$. In this limit case, continuation payoffs after every history are always approximately equal to the subsequently expected negotiation payoffs, since the periods until new negotiations are discounted. This means it is not incentive compatible to take the outside option if negotiation payoffs are strictly above the outside option payoff. Equation (17) describes the negotiation payoffs of a repeated game in which players do not have the possibility to choose the outside option so that the disagreement point is determined by the inside option payoffs π^{io}. If these negotiation payoffs r^{io} exceed for both players their outside option payoffs, the outside option can indeed be ignored; taking it cannot be incentivized.

Also for the case that for some player i the outside option payoff exceeds this negotiation payoff, i.e. $\pi_i^{oo} > r_i^{io}$, there cannot be any RNE in which player i can negotiate some surplus above her outside option payoff, since then taking the outside option would not be incentive compatible. Then either no positive surplus can be generated and in the RNE player i immediately takes the outside option, or no RNE exists. For the latter case the proof of Proposition 4 characterizes the unique weak RNE payoffs as defined in Appendix A, and shows that player i then also only gets her outside option payoff.

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

5.3 Property rights and repeated hold-up

In this application we illustrate how the model of Section 2 can be extended to study the optimal allocation of property rights, which is a typical question from the hold-up literature (Hart and Moore (1990)), for a long-term relationship. More precisely, we construct a repeated negotiation version of Halonen (2002), who compares the relational efficiency of different ownership structures.

The game starts in state x_0, in which each player i can choose an investment $x_i \geq 0$ at cost $c(x_i)$. The cost function is assumed to be smooth, strictly increasing and strictly convex with $c(0) = c'(0) = 0$. In the next period in state (x_1, x_2), the players can agree to trade and split a surplus $S(x_1, x_2) = x_1 + x_2$.

If there is no agreement to trade, payoffs are determined by property rights and the asset specificity. Under *joint ownership*, players cannot use any assets outside the relationship and their no-trade payoffs are zero. If player 1 has *single ownership* over her assets, she gets a no-trade payoff of λx_1. The parameter $\lambda \in [0, 1]$ is a measure of asset specificity and describes how well the assets can be utilized outside the relationship. Joint ownership is equivalent to single ownership with $\lambda = 0$. Following Halonen (2002), we assume for simplicity that player 2 never has single ownership over any assets and always gets a no-trade payoff of zero.

In a one shot version of this game, single ownership and a larger value of λ help to reduce the hold-up problem. Player 1 has stronger incentives for investments than under joint ownership as her investments also increase her no-trade payoff and bargaining position. Yet, Halonen (2002) shows that if the two stage game is indefinitely repeated, joint ownership dominates single ownership for a sufficiently convex cost function because the lower no-trade payoffs acts as a more severe punishment in future interactions.

We study RNE in this infinite horizon setting and assume that after two periods, the game always starts again in state x_0. To be close to Halonen (2002), we assume that the players have equal bargaining power and that discounting with discount factor δ takes place only after a two-period cycle is finished.[18]

New negotiations can take place every period with negotiation probability ρ. We denote by

$$G(x) = S(x_1, x_2) - c(x_1) - c(x_2)$$

the joint payoff that is generated in a cycle with investments $x = (x_1, x_2)$ and trade. First best investments x^o maximize $G(x)$. Let

$$\bar{u}_1(\lambda) = \max_{x_1} \lambda x_1 - c(x_1)$$

[18]This is equivalent to a regular stochastic game with discounting after every period with a discount factor $\hat{\delta} = \sqrt{\delta}$ and the trade surplus and the no-trade payoffs multiplied by $1/\hat{\delta}$.

denote the maximum payoff player 1 can achieve without trade. We assume that players can already decide in x_0 to work alone and then receive the no-trade payoffs of $(\bar{u}_1, 0)$ for this cycle without any opportunity for trade, even if new negotiations take place in the 2nd period of the cycle. Like Halonen (2002) we study the conditions under which first best investments can be implemented.

Proposition 5. *An RNE with first best investments x^o exists if and only if*

$$\begin{aligned}\max_{x_1}\left(\rho\frac{1}{2}(S(x_1,x_2^o)+\lambda x_1)+(1-\rho)\lambda x_1-c(x_1)\right)+ \\ \max_{x_2}\left(\rho\frac{1}{2}(S(x_1^o,x_2)-x_1^o\lambda)-c(x_2)\right)-G(x^o) \quad \leq \\ \frac{\tilde{\delta}(1-\rho)}{(1-\tilde{\delta}(1-\rho))}\left(G(x^o)-\bar{u}_1(\lambda)\right)\end{aligned} \tag{18}$$

The first part of the left hand side of (18) is the sum of both players' expected payoffs from their best deviation in the one shot game. If no negotiations take place at the beginning of period 2, the deviating player is punished and only gets her no trade-payoff, otherwise the surplus is split according to the Nash bargaining solution. Moreover, the joint equilibrium path payoff $G(x^o)$ is subtracted from this term, so that the complete left hand side measures the joint incentives to deviate in the one shot game. The right hand side measures the scope of relational incentives by using punishment in future periods to prevent such deviations. It is the difference between the joint equilibrium path payoff and the players' no-trade payoff appropriately discounted and adjusted for the probability of new negotiations.

Without repeated negotiations ($\rho = 0$), the incentive constraint (18) can be simplified to $G(x^o) \geq \bar{u}_1(\lambda)$ and is always satisfied. As in the example from Section 2, any deviation from first best investments can be immediately punished by holding back in period 2 any trade surplus from the deviating player.

In general, the allocation of property rights generates similar trade-offs in our model as in Halonen (2002). A larger value of λ makes higher investments more attractive for player 1, as it increases the positive impact of investments on her bargaining position. On the other hand, a larger value of λ reduces the scope to punish today's deviations in future cycles. We can establish for RNE that the first effect always dominates the second.

Proposition 6. *Larger levels of λ can implement first-best investments in an RNE for a weakly larger set of δ and ρ, i.e, single ownership generally facilitates implementation of first best investments compared to joint ownership.*

This means that different from Halonen (2002), comparative statics with respect to ownership structure remain qualitatively the same as in the one-shot game. Halonen (2002) assumes that after any deviation from prescribed investments, the remaining surplus will be split equally and punishment consists of an indefinite repetition of the one-shot hold-up solution. A crucial difference of RNE is that punishment is not assumed to last forever but will also be newly negotiated with probability ρ each period. This reduces the importance of future punishment and optimal property rights are more strongly shaped by the familiar effects from static hold-up models.

5.4 Arms race

Arms races have been widely studied with economic models and game theory. One branch of literature, including Jervis (1978), Baliga and Sjöström (2004) and Abbink, Dong and Huang (2020), models arms races as coordination failure assuming that if a country could be sure that other countries do not invest into arms, it also has no incentives for such investment. Another branch, including van der Ploeg and de Zeeuw (1990) and Garfinkel (1990), study macroeconomic models that include an explicit term in a country's utility function that increases in the own weapon arsenal and decreases in the other country's weapons arsenal. Garfinkel (1990) refers to this function as a *tribute* function with the idea that a larger weapon arsenal can be used to extract payments or favors from other countries. This provides a motive for arms races beyond coordination failure.

Our example explores the microfoundation and implications of such a tribute motive in a two-player setting where the costs of attacks exceed any direct benefits. In Pareto optimal SPE and in MPE arms races do not occur, but they naturally arise in our model of repeated negotiations. We denote by $x_i \in \{0, 1, ..., \bar{x}\}$ the integer valued weapon arsenal of player $i = 1, 2$. A state $x = (x_1, x_2)$ describes the arsenals of both players. A weapon arsenal x_i causes maintenance costs of $c_m x_i$ each period for player i. As long as x_i is below $\bar{x}$ a player can invest in order to attempt to increase his weapon arsenal by one unit. An investment successfully increases the arsenal only with probability π_I and involves costs of c_I. A player can also decrease its arsenal by one unit at no cost.

Each period each player also decides whether or not to attack the other player. If player i attacks, it creates harm that reduces the other player's payoff by x_i, while player i bears attack costs of cx_i. All cost coefficients are strictly positive.

Fact 1. *In every Pareto optimal SPE and in the unique Markov perfect equilibrium no weapons are bought or used.*

There is no direct benefit from acquiring, maintaining, and using weapons: it only involves costs. The only reason to acquire weapons can be that the threat of using them increases a player's bargaining position and allows him to extort payments from the other player. However, in a Pareto optimal SPE players can perfectly coordinate to ignore such threats and doing so is welfare optimal. In an MPE attacks have no consequences for future periods since they do not change the state. Attacks will thus never be performed due to their costs. Hence, no player

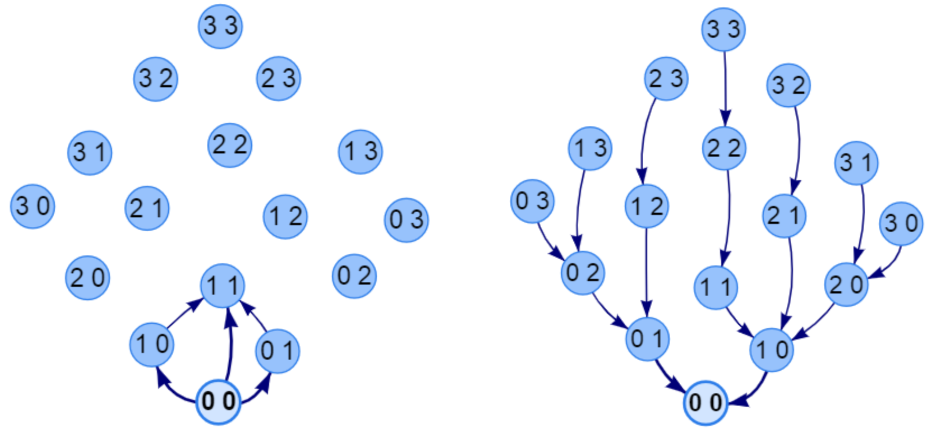

Figure 4: T-RNE of two arms race games with $\bar{x} = 3$, $\delta = 0.99$, $\rho = 0.65$, $c_I = 0.01$, $c = 0.05$, $c_m = 0.3$, $T = 1000$. The left panel shows the case of a small success probability of investments ($\pi_I = 0.08$), and the right hand panel the case that investments are always successful ($\pi_I = 1$). A directed arrow indicates a positive transition probability on the equilibrium path from one state to another. No arrow is drawn to indicate a positive probability to remain in the current state.

has an incentive to bear the costs of acquiring weapons. For a similar reason, no weapons are acquired in an RNE if one assumes that disagreement payoffs correspond to MPE payoffs of the truncated game.

However, assuming optimal punishments under disagreement, arms races can occur in RNE for intermediate negotiation probabilities. Figure 4 illustrates the state transitions for two numerically solved T-RNE in which negotiations only take place until period $T = 1000$ (recall Section 4.5).

The left panel corresponds to a game in which weapon investments only have a small success probability of 8%. We see that in the state $(0, 0)$ without weapons both players try to increase their arsenal. If both investments have been successful, play moves to state $(1, 1)$ and then no more investments take place. Players can then use the acquired weapons to punish any further investment attempts and thereby effectively prevent a continuation of the arms race.

If only one player successfully acquired weapons (state $(0, 1)$ or state $(1, 0)$) then only the other player continues investing until state $(1, 1)$ is reached. This result may seem surprising: one might have suspected that e.g. in the state $(1, 0)$ player 1 can use the threat of an attack to prevent player 2 from investing, while player 1 continues investing himself. To understand the result, first note that player 1 can extort higher payments from player 2 in state $(1, 0)$ on the equilibrium path if he allows player 2 to invest. Preventing investments by player 2 would require that player 2 has to make higher payments (backed by an attack threat) when deviating by investing than on the equilibrium path. But that means that the extorted payments on the equilibrium path cannot not be too high. In the illustrated equilibrium player 1 thus reaps relatively high payoffs while being in state $(1,0)$ but accepts that after some time player 2 will have successfully invested and then no payments can be extorted anymore.

There are indeed multiple T-RNE for the given parametrization. In another T-RNE the state $(1,0)$ transits to state $(2,0)$ instead. In that equilibrium player 1 extracts smaller payments in state $(1,0)$ but can extort higher payments in the future once the state $(2,0)$ is reached. Given that all T-RNE have the same unique negotiation payoffs, we can rule out, however, T-RNE in which starting from state $(0,0)$ players ever permanently end up in a state in which the total weapon arsenal is different from 2. Since maintenance costs are relatively high in our setting, a major priority of negotiations is to keep the total weapon arsenal as small as can possibly be incentivized, and the total punishment potential from a total arsenal size of 2 suffices to achieve this goal.

The pane on the right in Figure 4 illustrates state transitions in a T-RNE assuming that investments are now always successful ($\pi_I = 1$) while keeping all other parameters the same as in the previous scenario. One might suspect that this change yields more weapon investments on the equilibrium path since a higher success probability without changing the investment costs effectively makes investments cheaper. Yet, we actually find that now in no state weapon investments are conducted. Moreover, if players start with strictly positive weapon arsenals, they reduce their weapon arsenals step by step until the weapon-free state $(0,0)$ is reached.

To understand the result, note that players invest as punishment (and attack if the current arsenal of the punisher is positive). Since investments are always successful, the threat of a quick acquisition of weapons, which will be used if the punished player refuses to pay a fine, is sufficient to prevent any weapon acquisition on the equilibrium path. Moreover, this threat is so effective that it can even incentivize players to reduce their weapon arsenal. This outcome is facilitated by the fact that a large weapon arsenal generally gives less scope for extortion if the other player can quickly and cheaply build his own arsenal.

One can easily show that for all parametrizations of the game above attacks will never take place on the equilibrium path but will only be used as an off-equilibrium path threat. There is simply no gain for any player to perform attacks on the equilibrium path. However, consider now a variant of the game in which an attack by player i destroys with probability $\phi x_i/\bar{x}$ one unit of the other player's weapons arsenal. Figure 5 shows the corresponding T-RNE for the case $\phi = 0.5$, a small success probability of investments $\pi_I = 0.08$, and all other parameters as in the previous scenarios.

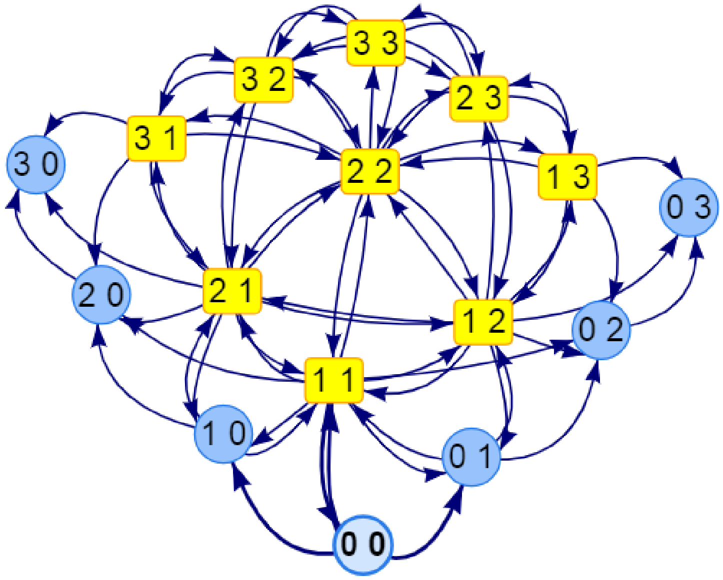

Figure 5: T-RNE of the arms race where an attack can destroy other player's weapons and small success probability of investments. Blue circles correspond to states without attack on equilibrium path, yellow boxes to states with attacks on the equilibrium path.

We now find that in many states attacks take place on the equilibrium path. The rationale for an equilibrium path attack is that it can directly reduce the other player's weapon arsenal and thus put the attacker into a better future bargaining position. Even though players negotiate and would try to prevent such wasteful actions, the incentives for attacks are simply too large for attacks to be avoided. We see that in this variant of the game also investments take place in almost all states and that there are only two absorbing states: the extremes $(3,0)$ and $(0,3)$ in which one player has obtained a dominant position.

6 Concluding remarks

This paper has been motivated by the discrepancy between the behavioral assumptions of relational contracting models and hold-up models. An important facet of this discrepancy is how and to what extent players will make themselves vulnerable in repeated relationships. Pareto efficient PPE predict that players not only do not take explicit measures against being held-up, but are willing to actively make themselves vulnerable, without accounting for any risk of a weaker bargaining position in future interactions. We have studied how repeated negotiation equilibria reconcile relational contracting with hold-up concerns in a tractable fashion by assuming that relational contracts are repeatedly newly negotiated. In our model, the advantage of creating vulnerabilities through formal contracts or institutions is traded off against the effects on future bargaining positions. We have shown that players prefer to increase their vulnerability gradually rather than committing to a high level from the start of the relationship.

With the probability of new negotiations ρ we have introduced a parameter that measures the importance of history-independent bargaining power. Factors that could affect the negotiation probability in applications include the frequency of communication, how well agreements are recorded, technological progress, the competitiveness of the environment, the duration of the relationship, or the frequency of changes of the external environment. Studying these applications in detail to derive new testable predictions is an area of future research.

A related straightforward extension of our model are negotiation probabilities $\rho(x)$ that depend on the current state.[19] One application are models in which negotiation probabilities decrease over time or after streaks of successful cooperation, formalizing the idea that longer relationships become more stable over time. State-dependent negotiation probabilities could also be used to model incompleteness of relational contracts. To model that some state is not considered in an initial relational contract, one can assign a negotiation probability of one to such states while having smaller or zero negotiation probabilities for states that have been initially considered. A further application of state-dependent negotiation probabilities is to model different forms of endogenous negotiations, by letting players choose to move to states in which new negotiations take place.

[19]The existence and uniqueness result in strongly directional games (Proposition 3) and the general existence of weak RNE (Theorem 1 in Appendix A) also hold for this more general model.

Appendix A: Weak RNE and Existence

Another example of non-existence of RNE: The blackmailing game

In this appendix, we introduce the concept of weak RNE, for which we prove a general existence result. For better intuition, we first provide another example where RNE fail to exist: the blackmailing game. 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 moves to an absorbing state x_1, in which no actions can be taken. Stage game payoffs are

$$\pi(x_0, a_S) = (0, 1)$$
$$\pi(x_1) = \pi(x_0, a_R) = (0, 0).$$

Revealing the evidence involves no cost for the blackmailer but reduces the target's payoffs by 1 in the current and all future periods. We assume that disagreement payoffs correspond to the worst continuation payoffs in the corresponding truncated games. Simple arguments show that no RNE exists in this game for $0 < \rho < 1$. One might suspect that by the threat to reveal the evidence the blackmailer could extract money from the target and thereby get a positive negotiation payoff in the initial state, i.e., $r_1(x_0) > 0$. Yet, this is not possible for a positive negotiation probability. A truncated game $\Gamma(x_0, r)$ with $r_1(x_0) > 0$ has no SPE in which the blackmailer reveals the evidence. The reason is that in state x_1 payoffs are 0 forever. Hence, the blackmailer would always prefer to stay in state x_0, predicting that he will again extract money from the target in future negotiations.

If instead the blackmailer has zero negotiation payoffs in state x_0, it becomes credible to reveal the information in the corresponding truncated game. Both players would then have punishment payoffs of zero in the truncated game. The Nash bargaining solution then implies that the blackmailer gets a payoff equal to his bargaining weight, since

$$r_1(x_0) = \bar{v}_1 + \beta_1(\bar{U}(x_0) - \bar{v}_1 - \bar{v}_2) = \beta_1. \tag{19}$$

Hence, if the blackmailer has a strictly positive bargaining weight, this contradicts the requirement that the blackmailer must get zero negotiation payoffs. Therefore, no RNE exists.

Weak repeated negotiation equilibrium

The reason why RNE may not exist is that $\mathcal{U}(x, r)$ is in general not lower hemi-continuous in the parameter r and therefore the mapping from future to current negotiations payoffs may not have a fixed point. The definition of weak RNE therefore relies on *stable* equilibrium payoffs, which are robust to perturbations in r. We say that a profile of PPE payoffs u in the truncated game $\Gamma(x, r)$ is *stable* if the correspondence $\tau \mapsto \mathcal{U}(x, \tau)$ has, locally around r, a continuous selection through the point (r, u).

Definition 3. A PPE payoff profile $u \in \mathcal{U}(x, r)$ is *stable* if there exists a neighborhood $\mathcal{N}$ of r and a continuous function $f : \mathcal{N} \to \mathbb{R}^n$ such that $f(r) = u$ and $f(\tau) \in \mathcal{U}(x, \tau)$ for all $\tau \in \mathcal{N}$.

For example, a PPE payoff profile $u \in \mathcal{U}(x, r)$ is stable if it is the payoff of a PPE σ that stays a PPE for small changes in r. Stability is a generic property in the sense that it holds for a dense set of negotiation payoffs. However, recall the truncated game in the blackmailing game in which the blackmailer has zero negotiation payoff. PPE payoffs of $\Gamma(x_0, 0)$ that grant the proposer a positive payoff are not stable, because they cease to be equilibrium payoffs if the blackmailer's negotiation payoff becomes slightly positive. The key assumption of a weak RNE is that such unstable continuation payoffs can be ignored in negotiations. Let

$$\mathcal{U}_d(x, r) = \{u \in \mathcal{U}(x, r), u_i \geq d_i(x, r) \text{ for all } i = 1, ..., n\} \tag{20}$$

denote the set of PPE payoffs that grant every player i at least the payoff $d_i(x, r)$.

Definition 4. For given bargaining weights β and disagreement point function d, a PPE $\sigma \in \Sigma^*$ is a *weak repeated negotiation equilibrium* if for all states $x \in X$ its negotiation payoffs r^σ can be written as $r^\sigma(x) = \sum_{i=1}^n \beta_i r(x, i)$ such that $r(x, i) \in \mathcal{U}_d(x, r)$ and $r_i(x, i) \geq \tilde{u}_i$ for all stable PPE payoffs $\tilde{u} \in \mathcal{U}_d(x, r)$.

The idea behind this definition is that in the negotiation phase, "nature" determines whether new negotiations are initiated and by whom. With probability β_i, the continuation equilibrium is selected by player i, who picks his preferred continuation equilibrium under the constraint that the other players have to receive at least their disagreement payoff. An RNE would correspond to the case that each proposer picks the continuation equilibrium that gives him the maximum payoff and the other players their disagreement payoffs. In contrast, in a weak RNE, the proposer only compares his payoff to the stable payoffs. Hence, as the name suggests, weak RNE is indeed a weaker concept.

Note that the characterization result in Proposition 1 continues to hold: For every weak RNE there exists a simple weak RNE with the same negotiation payoffs. To guarantee existence of a weak RNE, we have to make a stronger assumption about disagreement payoffs.

Assumption 2. *Let $d(x, r) \in \mathcal{U}(x, r)$ be either equal to the lowest possible payoff for each player, $d(x, r) = \bar{v}(x, r)$, or a continuous function in r.*

While this assumption already implies that $\mathcal{U}(x, r)$ is nonempty, $\mathcal{U}(x, r)$ is generally nonempty if one assumes that action sets are finite and mixed actions are allowed (see e.g. Sobel (1971) for existence of Markov equilibria in stochastic games).

Theorem 1. *If Assumption 2 holds, a weak RNE exists.*

In general, there can be multiple weak RNE, as we will illustrate in the examples below. We propose to select a weak RNE that is not Pareto dominated by any other weak RNE.

Weak RNE in the blackmailing game

The blackmailing game has a weak RNE with negotiation payoff $r(x_0) = (0, 1)$. Although the blackmailer's Nash bargaining payoff would be larger for these negotiation payoffs, these payoffs are not stable, since the threat to reveal the evidence is credible only if $r_1(x_0) = 0$. This also implies that there can be no weak RNE in which the blackmailer has a positive negotiation payoff. There are, however, weak RNE in which the blackmailer reveals the evidence with positive probability on the equilibrium path. This means that there are multiple weak RNE payoffs that differ only in player 2's payoff, ranging from her Nash bargaining payoff up to a payoff of one. When restricting attention to weak RNE that are not Pareto dominated by other weak RNE, the unique payoff profile is $(0, 1)$.

Intuitively, one can interpret this weak RNE outcome as follows. Assume player 2 promises the blackmailer to pay an amount ε each period for not revealing the evidence. Since any positive payment destroys the blackmailer's incentive to reveal the evidence, the weak RNE outcome corresponds to the limit case of $\varepsilon \to 0$. By a similar intuition the T-RNE payoff profiles also converge to $(0, 1)$ as T goes to infinity.

Extension: brinkmanship

The blackmailer may be able to extort positive payments if there is the possibility to conduct brinkmanship (Schelling (1960) and Schwarz and Sonin (2008)).[20] 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 when 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 brinkmanship action that reveals the evidence only with probability $\phi \in (0, 1)$.

Proposition 7. *For all $\phi \in (0, 1)$, there exists a weak RNE in which the brinkmanship action is used by the blackmailer to extort a payment from the target. This payment is maximized if*

[20] 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 (2008) show that such a divisible threat can dramatically increase the bargaining value of an otherwise non-credible threat by making punishment possible in a subgame-perfect equilibrium. In contrast, in our model revealing the evidence can be part of a subgame perfect equilibrium and the commitment problem instead results from renegotiation.

ϕ is equal to $\phi^* = \frac{(1-\delta)(1-\rho)}{\rho\beta_1+(1-\delta)(1-\rho)}$. For $\phi \leq \phi^*$ the weak RNE is also an RNE and for $\phi > \phi^*$ only weak RNE but no RNE exist.

A larger value of ϕ means a harsher punishment which can be used to extract larger payments. Therefore, the blackmailer's negotiation payoff increases in ϕ as long as ϕ is small enough. However, once ϕ exceeds ϕ^*, credibility of the punishment imposes an upper bound on the blackmailer's negotiation payoff. An RNE then fails to exist and the blackmailer's negotiation payoff in the weak RNE decreases in ϕ. Note that ϕ^* decreases in the blackmailer's bargaining power β_1, the frequency of renegotiation ρ, and the weight on future payoffs δ. Hence, the optimal brinkmanship action has a lower probability of revealing the evidence, meaning that punishment needs to be more gradual, the more severe the blackmailer's commitment problem is.

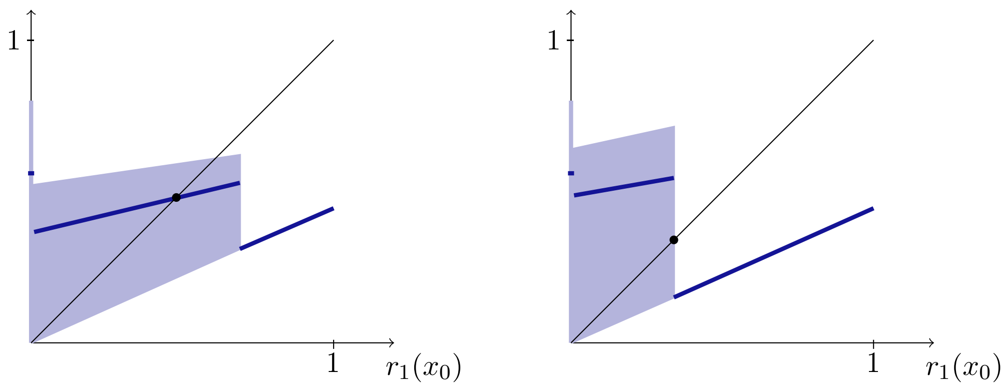

Figure 6: This figure shows the blackmailer's payoff set $\mathcal{U}_1(x_0, r)$ (light blue) and Nash bargaining payoff (dark blue) for each possible negotiation payoff $r_1(x_0)$ for the parameters $\rho = 0.7$, $\delta = 0.8$, $\beta_1 = 0.7$ and $\phi = 0.4$ (left), respectively, $\phi = 0.7$ (right).

Figure 6 provides some graphical intuition for the fixed point conditions determining an RNE (left) and a weak RNE (right). The shaded area shows the set of player 1's SPE payoffs in the truncated game in state x_0 in dependence on $r_1(x_0)$, assuming $r_2(x_0) = 1 - r_1(x_0)$. As a fixed point, every (weak) RNE must lie on the 45° line. An RNE (left panel) also lies on the blue line that shows player 1's bargaining payoffs in the corresponding truncated games. In the case that only a weak RNE exists (right panel), there is no such intersection, instead player 1 gets the highest negotiation payoff for which it is still credible to use the brinkmanship action. One can interpret the bargaining outcome in this weak RNE as follows: It is not possible that player 1 can negotiate an even higher transfer by player 2 since both players would expect that such a negotiation outcome would then be repeated in future negotiations. This mutual expectation would destroy the incentive compatibility of player 1's threat to punish via brinkmanship. Player 2 has therefore no reason to accept higher payment demands in negotiations.

Weak RNE in Gangsta's Paradise game

Consider now the Gangsta's Paradise game with a parameter constellation for which no RNE exists. These are cases in which mutual cooperation cannot be implemented in state x_0 if future negotiations in x_0 implement mutual cooperation (without money burning), but mutual cooperation could be implemented if future negotiations in x_0 implement mutual defection. Here the weak RNE outcome involves money burning in state x_0 together with mutual cooperation. Money burning must reduce both players' negotiation payoff so much that staying in state x_0 becomes sufficiently unattractive. More precisely, negotiation payoffs $r_i(x_0)$ must decrease so far that incentive constraint (12) holds with equality. Note that it is not crucial that players literally burn money, just that the negotiation outcome entails a particular inefficiency. Money burning could for example be replaced by using a public correlation device to coordinate on playing mutual defection with a certain probability.

First period negotiations and weak RNE

The intuition for money burning in the negotiation outcome in the Gangsta's Paradise game seems more plausible if players negotiate after period 1 in state x_0, after some player has previously defected. Yet, one may argue that there is no need for money burning in the first period, because negotiations in the first period are likely perceived differently from negotiations when state x_0 is reached after a defection. One can model this assumption by simply assuming that in the first period the game starts in a separate state x_F. Then a weak RNE exists that implements mutual cooperation without money burning in the first period and would only entail money burning if players negotiate in state x_0 in the future. While also weak RNE with money burning in the first period would still exist, they are Pareto dominated by the weak RNE without money burning.

Similarly, we also find for the blackmailing game weak RNE with positive payoffs for the blackmailer if one assumes that the first period is a separate state. Importantly, however, such first period effects do not arise in RNE. It is straightforward to show that the set of RNE payoffs does not change when relabelling the initial state in the first period as a separate state.

Appendix B: Algorithm to compute RNE in strongly directional games

Step 1

We first compute for each absorbing state x the set of pure SPE payoffs of the repeated game given a discount factor $\tilde{\delta} = (1-\rho)\delta$, e.g. using the algorithm described in Goldlücke and Kranz (2012). With the maximum joint payoffs $\tilde{U}(x)$ and punishment payoffs $\tilde{v}_i(x)$, we can compute the negotiation payoffs $r_i(x) = \tilde{d}_i(x) + \beta_i(\tilde{U}(x) - \sum_{j=1}^n \tilde{d}_j(x))$. The payoff set of the truncated game starting in state x is then characterized by the maximum joint payoffs $\bar{U}(x) = \tilde{U}(x)$ and punishment payoffs $\bar{v}_i(x) = \frac{\delta\rho r_i(x)+(1-\delta)\tilde{v}_i(x)}{1-\tilde{\delta}}$. When this step is completed, all absorbing states are solved in the sense that negotiation payoffs and payoff set are known. We denote by X_s the set of solved states.

Step 2

We pick some unsolved state x from which the game can only transit to solved states. For all pure action profiles $a \in A(x)$ we compute the joint equilibrium payoffs as in (4),

$$U(x|a) = (1-\delta)\Pi(x,a) + \delta E[\bar{U}(x')|x,a],$$

and punishment payoffs for each player $i = 1, 2$ as in (5)

$$v_i(x|a) = \max_{\hat{a}_i \in A_i(x)} \left[(1-\delta)\pi_i(x, \hat{a}_i, a_{-i}) + \delta E[(1-\rho)\bar{v}_i(x') + \rho r_i(x')|x, \hat{a}_i, a_{-i}\right].$$

This allows us to determine the set of pure action profiles $\hat{A}(x)$ that can be implemented in state x as all action profiles $a \in A(x)$ that satisfy the summed incentive constraint

$$U(x|a) \geq v_1(x|a) + v_2(x|a).$$

The set of (pure) SPE payoffs of the truncated game starting in state x is then characterized by the joint payoff $\bar{U}(x) = \max_{a\in\hat{A}(x)} U(x|a)$ and for each player i the punishment payoffs $\bar{v}_i(x) = \min_{a\in\hat{A}(x)} v_i(x|a)$. Since negotiation payoffs of the solved states are known, renegotiation payoffs in state x can be easily computed as $r_i(x) = d_i(x,r) + \beta_i(\bar{U}(x) - \sum_{j=1}^n d_j(x,r)$. We mark state x as solved and repeat step 2 until all states are solved.

Appendix C: Proofs

Proof of Lemma 1

Consider a public strategy profile of the original game for which only the history since the last negotiations matters, $\sigma \in \Sigma^*$. Let h^0 be a history of a truncated game $\Gamma(x, r^\sigma)$, and let $h^{signal}(h^0)$ be the history h^0 but with negotiation signals $R_0 = 1$ and $R_t = 0$ for all $t \geq 1$. The tuple $(\sigma^x)_{x\in X}$ of public strategy profiles of the truncated games $\Gamma(x, r^\sigma)$ corresponding to σ is defined by $\sigma^x(h^0) = \sigma(h)$ for all histories h^0 of the truncated game, where h is a history of the original game with $h^C(h) = h^{signal}(h^0)$. Then $\sigma|_h$ in the original game and $\sigma^x|_{h^0}$ in the truncated game $\Gamma(x, r^\sigma)$ prescribe the same play until a new positive negotiation signal/an absorbing state is reached. By definition, continuation payoffs of σ following a positive negotiation signal are the same as the payoffs in the corresponding absorbing states. Therefore, a profitable one-shot deviation from σ at history h induces a profitable deviation from σ^x at h^0, and vice versa. $\square$

Proof of Proposition 1 (Simple RNE)

Let σ be a RNE with negotiation payoffs r^σ. Goldlücke and Kranz (2018) show that $\Gamma(x_0, r^\sigma)$ has an *optimal* simple PPE $\hat{\sigma}$ such that for any state $x \in X$, varying the up-front transfers of $\hat{\sigma}|_h$, where h is an arbitrary history ending in state x, yields all PPE payoffs of $\Gamma(x, r^\sigma)$. Let $\hat{\sigma}$ be such an optimal simple PPE in the truncated game $\Gamma(x_0, r^\sigma)$. We can construct a simple RNE $\sigma^* \in \Sigma^*$ from $\hat{\sigma}$ by defining transfers following a negotiation signal. Let $h^0 = (..., (x, 0))$ be a history that ends in state x with no new negotiations, and let $u(\hat{\sigma}|_{h^0}, x_0, r^\sigma)$ denote the continuation payoff created by $\hat{\sigma}$ in the truncated game $\Gamma(x_0, r^\sigma)$ following such a history. Define for all $x \in X$ the payments $p^0(x)$ such that $u(\hat{\sigma}|_{h^0}, x_0, r^\sigma) - p^0(x) = r^\sigma(x)$. Let $\sigma^*(h) = p^0(x)$ for $h^C(h) = (x, 1)$ and otherwise $\sigma^*(h) = \hat{\sigma}(h^c(h))$. Then σ^* is a simple PPE and has the same payoffs as σ following a positive negotiation signal and is therefore a weak repeated negotiation equilibrium. $\square$

Proof of Proposition 2 (Repeated Games)

For a given strategy profile σ, let $\tilde{\pi}(t, \sigma)$ denote the expected payoffs in period t in a repeated game with negotiation probability equal to zero. For fixed negotiation payoffs r, the expected payoff in the truncated game $\Gamma(r)$ can be written as

$$\begin{aligned} u(\sigma, r) &= \sum_{t=0}^{\infty} (\delta(1-\rho))^t \left((1-\delta)\tilde{\pi}(t,\sigma) + \delta\rho r\right) \\ &= (1-\delta) \sum_{t=0}^{\infty} (\delta(1-\rho))^t \, \tilde{\pi}(t,\sigma) + \frac{\delta\rho r}{1-(1-\rho)\delta}. \end{aligned} \tag{21}$$

This is a positive affine transformation of the payoff of the repeated game without negotiations with discount factor $\tilde{\delta} = \delta(1-\rho)$. Hence, the set of public perfect equilibria in the truncated game is independent of the negotiation payoffs r and the same as in the repeated game with discount factor $\tilde{\delta}$. Defining

$$a = \frac{1-\delta}{1-\tilde{\delta}},$$

we have that

$$\frac{\delta\rho r}{1-(1-\rho)\delta} = (1-a)r.$$

The PPE payoff set of the truncated game $\Gamma(r)$ is equal to

$$\mathcal{U}(r) = \{au + (1-a)r : \; u \in \mathcal{U}(\tilde{\delta})\}.$$

Hence, the maximum attainable joint surplus in the truncated game $\Gamma(r)$ is equal to

$$U(r) = a\bar{U}(\tilde{\delta}) + (1-a)\sum_{i=1}^{n} r_i, \tag{22}$$

and the minimum attainable payoff for player i is equal to

$$\bar{v}_i(r) = a\bar{v}_i(\tilde{\delta}) + (1-a)r_i.$$

Given Assumption 1, it holds that $d(r) \in \mathcal{U}(r)$. Moreover, since for any strategy profile σ we have that $u(\sigma, r) = u(\sigma, 0) + (1-a)r$ it must also be true that

$$d(r) = d(0) + (1-a)r. \tag{23}$$

An RNE exists if the equation

$$r_i = d_i(r) + \beta_i(U(r) - \sum_{j=1}^{n} d_j(r)) \tag{24}$$

is satisfied for some negotiation payoff r. Using (23) and (22), this equation can be rewritten as

$$r_i = \frac{1}{a}d_i(0) + \beta_i(U(\tilde{\delta}) - \sum_{j=1}^{n} \frac{1}{a}d_j(0)). \tag{25}$$

This shows that an RNE exists with these payoffs. If the disagreement points are always the minimum PPE payoffs, then $d_i(0) = \bar{v}_i(0) = a\bar{v}_i(\tilde{\delta})$. If the disagreement payoffs result from playing a Nash equilibrium α^{NE}, then $d_i(0) = a\pi_i(\alpha^{NE})$. □

Proof of Proposition 3 (Strongly Directional Games)

We can order the states such that for $l < m$ the state x_l cannot be reached from state x_m. There exists an $\bar{m}$ such that for all $m \geq \bar{m}$, the states x_m are absorbing states. For the absorbing states, the results for repeated games apply given that Assumption 1 is imposed. Hence, for all states x_m with $m \geq \bar{m}$ there is a unique RNE payoff $r(x_m)$. Let all negotiation payoffs in those states be fixed at $r(x_m)$. Consider $l < \bar{m}$ such that for all states x_m with $m > l$ there is such a unique fixed RNE payoff. Since state x_l is only visited once, all future negotiation payoffs are fixed. There is a unique largest joint payoff $\bar{U}(x_l)$ in the game that starts in this state. Hence, there is also a unique RNE payoff, which we can fix as the negotiation payoff in state x_l. Note that if a game is only weakly directional like the Gangsta's paradise game, in which players can stay in state x_0 or move to paradise, then changing equilibrium actions in state x_0 also changes future negotiation payoffs. Here, however, there is no interdependency between equilibrium actions in state x_l and future negotiation payoffs. By induction, we eventually find the unique RNE payoff of the initial state. □

Proof of Proposition 4 (Outside Option Principle)

This proof finds all weak RNE (see Appendix A) to then show that they all fulfill the proposition's claim. 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. Define $S(e) = \pi_1^{io} + \pi_2^{io} + e - k(e)$. Fix negotiation payoffs $r = r(x_0)$ and consider first the case that for both players the outside option payoff is lower than the minimum payoff that they could get in the relationship:

$$\pi_i^{oo} < \frac{(1-\delta)\pi_i^{io} + \delta\rho r_i}{1-\tilde{\delta}} \text{ for } i = 1, 2. \tag{26}$$

Since taking the outside option is not credible in this case, it holds that $v_i(x_0, r) = \frac{(1-\delta)\pi_i^{io}+\delta\rho r_i}{1-\tilde{\delta}}$ for both players. An effort level e can then be played on the equilibrium path of a simple equilibrium in the truncated game if and only if

$$\tilde{\delta}(S(e) - S(0)) \geq (1-\tilde{\delta})k(e). \tag{27}$$

As in a repeated game, this condition does not depend on the negotiation payoffs. Define e^{io} as the effort level that maximizes $S(e)$ among all e that satisfy (27). Since all payoffs are stable in this truncated game, if a weak RNE exists in this parameter region it must be an RNE with payoffs equal to $r_i^{io} = \pi_i^{io} + \beta_i(S(e^{io}) - S(0))$. An RNE with this payoff indeed exists if condition (26) is satisfied for these negotiation payoffs, possibly with equality[21], i.e., if

$$(\pi_i^{oo} - \pi_i^{io})(1-\tilde{\delta}) \leq \delta\rho\beta_i(S(e^{io}) - S(0)) \text{ for } i = 1, 2. \tag{28}$$

In particular, for $\rho > 0$ and $\delta \to 1$, condition (28) becomes $\pi_i^{oo} \leq r_i^{io}$. For $\rho = 0$, this type of RNE does not exist.

Next, consider the case that taking the outside option is credible for at least one player:

$$\pi_i^{oo} > \frac{(1-\delta)\pi_i^{io} + \delta\rho r_i}{1-\tilde{\delta}} \text{ for some } i \in \{1, 2\}. \tag{29}$$

In this case, $v_j(x^0, r) \approx \pi_j^{oo}$ for both $j = 1, 2$. An effort e can be sustained with the threat of the outside option as a punishment if

$$\frac{(1-\delta)S(e) + \delta\rho R}{1-\tilde{\delta}} \geq (1-\delta)(\pi_1^{io}+e) + \delta\rho r_1 + (1-\tilde{\delta})\max\{\pi_2^{oo}, \frac{(1-\delta)\pi_2^{io} + \delta\rho r_2}{1-\tilde{\delta}}\} + \tilde{\delta}(\pi_1^{oo}+\pi_2^{oo}).$$

[21] If (26) holds with equality, the truncated game has SPE in which the outside option is taken but no stable payoffs that are larger than a player's maximum payoff with inside option punishment.

Let e^{oo} be the effort level that maximizes $S(e)$ among all effort levels e that satisfy

$$\begin{aligned}
\tilde{\delta}(S(e)-(\pi_1^{oo}+\pi_2^{oo})) &\geq (1-\delta)(k(e)+\pi_2^{oo}-\pi_2^{io})-\delta\rho\beta_2(S(e)-(\pi_1^{oo}+\pi_2^{oo})) \\
\text{and } \tilde{\delta}(S(e)-(\pi_1^{oo}+\pi_2^{oo})) &\geq (1-\delta)k(e).
\end{aligned}$$

An RNE with surplus $\bar{U}=\max\{S(e^{oo}),\pi_1^{oo}+\pi_2^{oo}\}$ and negotiation payoffs

$$r_j^{oo}=\pi_j^{oo}+\beta_j(\bar{U}-(\pi_1^{oo}+\pi_2^{oo}))$$

indeed exists if (29) is satisfied for these negotiation payoffs, i.e., if

$$(1-\delta)(\pi_i^{oo}-\pi_i^{io})>\delta\rho\beta_i(\bar{U}-(\pi_1^{oo}+\pi_2^{oo})) \text{ for some } i\in\{1,2\}. \tag{30}$$

In particular, for $\rho=0$ an RNE with payoffs r^{oo} as described in the proposition exists. For $\rho>0$ and $\delta\to 1$, condition (30) becomes $\pi_1^{oo}+\pi_2^{oo}\geq\bar{U}$, which is the case that no effort can be sustained on the equilibrium path and $r_i=\pi_i^{oo}$.

If neither (28) nor (30) holds, then there must be a weak RNE that is not an RNE. The only case in which the truncated game has payoffs that are not stable is

$$\pi_i^{oo}=\frac{(1-\delta)\pi_i^{io}+\delta\rho r_i}{1-\tilde{\delta}} \text{ and } \pi_{-i}^{oo}<\frac{(1-\delta)\pi_{-i}^{io}+\delta\rho r_{-i}}{1-\tilde{\delta}} \text{ for some } i\in\{1,2\}. \tag{31}$$

In this case, a small increase in r_i will make the outside option punishment payoff unattainable. The maximum stable payoff for player i is $\frac{(1-\delta)(S(e^{io})-S(0))}{1-\tilde{\delta}}+\pi_i^{oo}$. A weak RNE with negotiation payoffs $r_i=\pi_i^{oo}+\frac{(1-\delta)(\pi_i^{oo}-\pi_i^{io})}{\delta\rho}$ and $r_{-i}=S(e^{oo})-r_i$ exists if for $r=\beta_1 r^1+\beta_2 r^2$ it holds that r_i^i is larger than the maximum stable payoff for player i, which is equivalent to the following condition:

$$(1-\tilde{\delta})(\pi_i^{oo}-\pi_i^{io})\geq\delta\rho\beta_i(S(e^{io})-S(0)). \tag{32}$$

In addition, condition (31) has to be satisfied:

$$\delta\rho(S(e^{oo})-(\pi_1^{oo}+\pi_2^{oo}))>(1-\delta)(\pi_1^{oo}+\pi_2^{oo}-S(0)).$$

Note that for $\rho=0$, this parameter region is empty, and for $\rho>0$ and $\delta\to 1$, the weak RNE exists if $\pi_i^{oo}\geq r_i^{io}$, with negotiation payoff $r_i=\pi_i^{oo}$. □

Proof of Proposition 5 (Property rights and RNE)

Let $\bar{u}_1=\bar{u}_1(\lambda)$ and $\bar{u}_2=0$ be the payoffs that the players can receive if they do not work together. In state x_0, working alone yields the minmax payoff $\bar{u}_i$ for one two-period cycle and is an equilibrium because each player can decide unilaterally not to enter the productive relationship with the other player. In the truncated game starting in state x_0, the optimal punishment payoff is therefore

$$\bar{v}_i(x_0, r) = \frac{\delta\rho(2-\rho)r_1(x_0) + (1-\delta)\bar{u}_i}{1-\tilde{\delta}(1-\rho)}.$$

We are interested here in the conditions under which an RNE exists that has first best investments and trade decision on the equilibrium path. An RNE with such an equilibrium policy must have

$$r_i(x_0) = \bar{u}_i + \frac{1}{2}(G(x^o) - (\bar{u}_1 + \bar{u}_2)). \tag{33}$$

In state (x_1, x_2), the optimal punishment is not to trade, which is also incentive compatible for all δ and ρ. Consequently,

$$\bar{v}_1((x_1, x_2), r) = (1-\delta)\lambda x_1 + \frac{\delta\rho(1+\tilde{\delta})r_1(x_0) + \tilde{\delta}(1-\delta)\bar{u}_1}{1-\tilde{\delta}(1-\rho)} \tag{34}$$

and

$$\bar{v}_2((x_1, x_2), r) = \frac{\delta\rho r_2(x_0)}{1-\tilde{\delta}(1-\rho)}. \tag{35}$$

Using (33), negotiation payoffs in state x can be calculated as in a strongly directional game:

$$r_1(x_1, x_2) = (1-\delta)\frac{1}{2}(S(x_1, x_2) + \lambda x_1)) + \delta\frac{1}{2}(G(x^o) + \bar{u}_1), \tag{36}$$

and

$$r_2(x_1, x_2) = (1-\delta)\frac{1}{2}(S(x_1, x_2) - \lambda x_1)) + \delta\frac{1}{2}(G(x^o) - \bar{u}_1). \tag{37}$$

Note that $r_1(x_1, x_2) + r_2(x_1, x_2)$ are equal to the maximum joint surplus in that state, as always in an RNE in a game with perfect monitoring. There is indeed an RNE with these payoffs if and only if (8) holds for $a^e(x_0) = x^o$:

$$\begin{aligned} G(x^o) \quad \geq \quad & \max_{x_1}((1-\rho)\bar{v}_1((x_1, x_2^o), r) + \rho r_1(x_1, x_2^o) - (1-\delta)c(x_1)) \\ & + \max_{x_2}((1-\rho)\bar{v}_2((x_1^o, x_2), r) + \rho r_2(x_1^o, x_2) - (1-\delta)c(x_2)). \end{aligned}$$

Plugging in (34) and (35) yields

$$\begin{aligned} G(x^o)(1 - \frac{\tilde{\delta}\rho(1+\tilde{\delta})}{1-\tilde{\delta}(1-\rho)}) - \frac{(1-\rho)(1-\delta)\tilde{\delta}}{1-\tilde{\delta}(1-\rho)}\bar{u}_1 \quad \geq \quad & \max_{x_1}(\rho r_1(x_1, x_2^o) + (1-\delta)((1-\rho)\lambda x_1 - c(x_1))) \\ & + \max_{x_2}(\rho r_2(x_1^o, x_2) - (1-\delta)c(x_2)). \end{aligned}$$

Plugging in (36) and (37) yields

$$\frac{G(x^o)}{1-\delta}(1-\frac{\tilde{\delta}\rho(1+\tilde{\delta})}{1-\tilde{\delta}(1-\rho)}-\delta\rho)-\frac{(1-\rho)\tilde{\delta}}{1-\tilde{\delta}(1-\rho)}\bar{u}_1 \geq \max_{x_1}(\rho\frac{1}{2}(S(x_1,x_2^o)+\lambda x_1)+(1-\rho)\lambda x_1-c(x_1)) + \max_{x_2}(\rho\frac{1}{2}(S(x_1^o,x_2)-\lambda x_1^o)-c(x_2)).$$

Rearranging yields (18). □

Proof of Proposition 6 (Comparison of ownership structures)

Let $x^*(\lambda) = \arg\max_{x_1} \lambda x_1 - c(x_1)$ be the investment choice that maximizes player 1's payoff if she takes her assets outside of the relationship. Since the surplus $S(x_1, x_2) = x_1 + x_2$ is additively separable, it holds that $x^*(1) = x^o$ and $\bar{u}_1(1) = \frac{1}{2}G(x^o)$. Note also that $x^*(\lambda)$ is increasing in λ. By defining

$$F(\lambda) = G(x^o) + \frac{(1-\rho)\tilde{\delta}}{1-\tilde{\delta}(1-\rho)}(G(x^o)-\bar{u}_1(\lambda)) - \rho\frac{1}{2}(x_1^o+x_2^o) + \rho\frac{1}{2}\lambda x_1^o - \max_{x_1}(\rho\frac{1}{2}(x_1-\lambda x_1)+\lambda x_1-c(x_1)) - \max_{x_2}(\rho\frac{1}{2}x_2-c(x_2)),$$

we write condition (18) as $F(\lambda) \geq 0$. For $\lambda = 0$, which corresponds to joint ownership, it holds that

$$F(0) = G(x^o) + \frac{(1-\rho)\tilde{\delta}}{1-\tilde{\delta}(1-\rho)}G(x^o) - \rho\frac{1}{2}(x_1^o+x_2^o) - \max_{x_1}(\rho\frac{1}{2}x_1-c(x_1)) - \max_{x_2}(\rho\frac{1}{2}x_2-c(x_2)).$$

For $\lambda = 1$, it holds that

$$F(1) = \frac{1}{2}G(x^o) + \frac{\tilde{\delta}(1-\rho)}{1-\tilde{\delta}(1-\rho)}\frac{1}{2}G(x^o) - \rho\frac{1}{2}x_2^o - \max_{x_2}(\rho\frac{1}{2}x_2-c(x_2)) = \frac{1}{2}F(0).$$

Hence, if the first best can be implemented for $\lambda = 0$, then also for $\lambda = 1$. To show this also for $\lambda \in (0,1)$, we show that F is concave on this range, which implies that $F(\lambda) \geq F(1)$ for all λ. Applying the envelope theorem, the derivative of F is

$$F'(\lambda) = -\frac{\tilde{\delta}(1-\rho)}{1-\tilde{\delta}(1-\rho)}x_1^*(\lambda) + \rho\frac{1}{2}x_1^*(1) - (1-\rho\frac{1}{2})x_1^*(\rho\frac{1}{2}(1-\lambda)+\lambda).$$

Since x^* is increasing in λ, this derivative is indeed decreasing in λ. Hence, 1-ownership can implement efficiency for a weakly larger range of parameters than joint ownership. □

Proof of Theorem 1 (Existence Weak RNE)

Let $F(x)$ denote the convex and compact set of feasible payoffs of the game starting in state x and $\mathcal{R} = \Pi_x F(x)$. For all $r \in \mathcal{R}$, we define the correspondence $\phi : \mathcal{R} \to \mathcal{R}$ by

$$\phi(r) = \left\{ \hat{r} \in \mathcal{R} : \hat{r}(x) = \sum_{i=1}^{n} \beta_i \hat{r}^{x,i} \text{ s.t. } \hat{r}^{x,i} \in \mathcal{U}_d(x,r) \text{ and } \hat{r}_i^{x,i} \geq u_i \text{ for all stable } u \in \mathcal{U}_d(x,r) \right\}.$$

We first show that for a given $r \in \mathcal{R}$, the set $\phi(r)$ is a nonempty, compact and convex subset of $\mathcal{R}$. Since $d(x,r) \in \mathcal{U}(x,r)$, the set $\mathcal{U}_d(x,r)$ is nonempty. It is also a convex and compact subset of $F(x)$. Since the stable payoffs are a subset, there must exist an $\hat{r}^{x,i} \in \mathcal{U}_d(x,r)$ with $\hat{r}_i^{x,i} \geq u_i$ for all stable $u \in \mathcal{U}_d(x,r)$. Moreover, these inequalities remain true in the limit of a sequence of payoffs or for a convex combination of payoffs.

It remains to show that the correspondence ϕ is upper hemi-continuous to conclude that it has a fixed point. Let $r_m \to r$ and $\hat{r}_m \in \phi(r_m)$ with limit $\hat{r} = \lim_{m\to\infty} \hat{r}_m$. For all m, there are $\hat{r}_m^{x,i} \in \mathcal{U}_d(x,r_m)$ with $\hat{r}_m(x) = \sum_{i=1}^n \beta_i \hat{r}_m^{x,i}$ and $\hat{r}_{m,i}^{x,i} \geq u_i$ for all stable $u \in \mathcal{U}_d(x,r_m)$. Since $\mathcal{U}(x,\tau)$ is upper hemi-continuous in τ it must hold that $\hat{r}_m^{x,i} \to \hat{r}^{x,i} \in \mathcal{U}(x,r)$ with $\hat{r}(x) = \sum_{i=1}^n \beta_i \hat{r}^{x,i}$. According to Assumption 2, $d(x,r)$ is either continuous in r or equal to $\bar{v}(x,r)$. In both cases, $\tau \mapsto d_i(x,\tau)$ is lower semi-continuous. Therefore, $\hat{r}_{m,j}^{x,i} \geq d_j(x,r_m)$ implies $\hat{r}_j^{x,i} \geq d_j(x,r)$.

If for r there exists no stable $u \in \mathcal{U}_d(x,r)$, or none with $u_i > d_i(x,r)$, we already have that $\hat{r}_i^{x,i} \geq u_i$ for all stable $u \in \mathcal{U}_d(x,r)$. Otherwise, take any stable $\tilde{u}^{x,i} \in \mathcal{U}(x,r)$ with $\tilde{u}_i^{x,i} > d_i(x,r)$. By definition of stability, there exists a neighborhood N of r and a continuous function $f : N \to F(x)$ such that $f(r) = \tilde{u}^{x,i}$ and $f(\tilde{r}) \in \mathcal{U}(x,\tilde{r})$ for all $\tilde{r} \in N$. For m large enough such that $r_m \in N$ it must be true that $f(r_m) \in \mathcal{U}(x,r_m)$ is stable. To see this, note that to show that $f(r_m)$ is stable one can take a neighborhood N_m of r_m with $N_m \subset N$ and the same function f, restricted to N_m.

Assume first that $d_i(x,.) = \bar{v}_i(x,.)$. Since $f(r_m) \in \mathcal{U}(x,r_m)$ is stable, it must hold that $\hat{r}_{m,i}^{x,i} \geq f(r_m)_i$. In the limit $m \to \infty$ this becomes $\hat{r}_i^{x,i} \geq \tilde{u}_i^{x,i}$, and hence it holds that $\hat{r} \in \phi(r)$. Now assume that $d(x,\tau)$ is continuous in τ. Consider the sequence of payoffs $\tilde{u}_m$ with $\tilde{u}_{m,i} = f(r_m)_i$ and $\tilde{u}_{m,j} = d_j(x,r_m)$. Continuity of $d(x,.)$ implies that $d(x,r_m)$ is stable in $\mathcal{U}_d(x,r_m)$. For m large enough such that $f(r_m)$ is stable and $f(r_m)_i \geq d_i(x,r_m)$, the payoff vector $\tilde{u}_m$ is stable in $\mathcal{U}_d(x,r_m)$. For those m it must hold that $\hat{r}_{m,i}^{x,i} \geq f(r_m)_i$ and hence in the limit also here $\hat{r}_i^{x,i} \geq \tilde{u}_i^{x,i}$. We thus have shown that $\hat{r} \in \phi(r)$ in both cases.

Because of the Kakutani fixed point theorem, there must exist an $r \in \mathcal{R}$ with $r \in \phi(r)$, i.e., negotiation payoffs $r(x) = \sum_{i=1}^n \beta_i r^{x,i}$ with $r^{x,i} \in \mathcal{U}_d(x,r)$ and $r_i^{x,i} \geq u_i$ for all stable $u \in \mathcal{U}_d(x,r)$. There hence exists a tuple $(\sigma^x)_{x\in X}$ of PPE in the truncated games $\Gamma(x,r)$ with $u(\sigma,x,r) = r(x)$. Then the corresponding PPE $\sigma \in \Sigma^*$ (see Lemma 1) is a weak repeated negotiation equilibrium. $\square$

Proof of Proposition 7 (Brinkmanship)

A threat to punish player 2 (for not making a specified payment) with the brinkmanship action is easiest to implement if on the equilibrium path the evidence will never be revealed. This means the negotiation payoffs satisfy $r_1(x_0) + r_2(x_0) = 1$. A simple RNE can then have the brinkmanship action as punishment for player 2 if

$$(1-\phi) \geq v_1(x_0, r) + v_2(x_0, r), \tag{38}$$

with the punishment payoff of the target being

$$v_2(x_0, r) = \frac{(1-\phi)(1-\delta+\delta\rho r_2(x_0))}{1-(1-\phi)\delta(1-\rho)}$$

and the punishment payoff of the blackmailer

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

Condition (38) can thus be reformulated as

$$(1-\rho)(1-\delta(1-\rho))(1-\phi) \geq \rho r_1(x_0). \tag{39}$$

There exists an RNE if (39) is satisfied for the negotiation payoff given by $r_1(x_0) = v_1(x_0, r) + \beta_1(1 - v_1(x_0, r) - v_2(x_0, r))$, i.e. for the payoff

$$r_1(x_0) = r_1^{RNE}(x_0) \equiv \frac{\beta_1(1-\delta(1-\rho))\phi}{(1-\delta)(1-(1-\phi)\delta(1-\rho)) + \beta_1\delta\rho\phi}. \tag{40}$$

This is the case whenever

$$(1-\delta)(1-\rho)(1-\phi) \geq \rho\beta_1\phi, \tag{41}$$

which is equivalent to $\phi \leq \phi^*$. In this case, the blackmailer's payoff $r_1^{RNE}(x_0)$ is increasing in his bargaining weight β_1 and the revelation probability ϕ. Intuitively, larger revelation probabilities ϕ allow the blackmailer to extract larger payments in a subgame perfect equilibrium, and larger values of β_1 makes new negotiations more valuable for the blackmailer. If condition (41) does not hold for $r_1(x_0) = r_1^{RNE}(x_0)$, no RNE exist. Instead, there is a weak RNE in which the blackmailer gets the highest negotiation payoff for which the joint incentive constraint (39) is binding:

$$r_1(x_0) = r_1^{WRNE}(x_0) \equiv \frac{(1-\rho)}{\rho}(1-\delta(1-\rho))(1-\phi).$$

Note that the truncated game in state x_0 has SPE in which the blackmailer gets a higher payoff than $r_1^{WRNE}(x_0)$ by extracting even more money from player 2. These payoffs are not stable, because if the blackmailer would receive a future negotiation payoff slightly above $r_1^{WRNE}(x_0)$, the brinkmanship action could not be implemented anymore. $\square$

References

Abbink, Klaus, Lu Dong, and Lingbo Huang, "Arms Races and Conflict: Experimental Evidence," *The Economic Journal*, 08 2020, *131* (637), 1883–1904.

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.

Bagwell, Kyle and Robert W. Staiger, "A Theory of Managed Trade," *The American Economic Review*, 1990, *80* (4), 779–795.

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

Baliga, Sandeep and Robert Evans, "Renegotiation in repeated games with side-payments," *Games and Economic Behavior*, 2000, *33* (2), 159–176.

—— **and Tomas Sjöström**, "Arms races and negotiations," *The Review of Economic Studies*, 2004, *71* (2), 351–369.

Barrett, Scott, "The theory of international environmental agreements," *Handbook of environmental economics*, 2005, *3*, 1457–1516.

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," *The Quarterly journal of economics*, 1987, *102* (1), 147–159.

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.

____ **and** ____ , "Sunk costs and fairness in incomplete information bargaining," *Games and economic behavior*, 2005, *50* (2), 155–177.

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.

Garfinkel, Michelle R, "Arming as a strategic investment in a cooperative equilibrium," *The American Economic Review*, 1990, pp. 50–68.

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.

Gjertsen, Heidi, Theodore Groves, David A Miller, Eduard Niesten, Dale Squires, and Joel Watson, "Conservation Agreements: Relational Contracts with Endogenous Monitoring," *The Journal of Law, Economics, and Organization*, 2021, *37* (1), 1–40.

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

___ **and** ___ , "Renegotiation-proof relational contracts," *Games and Economic Behavior*, July 2013, *80*, 157–178.

___ **and** ___ , "Discounted stochastic games with voluntary transfers," *Economic Theory*, 2018, *66* (1), 235–263.

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.

Hall, Robert E and Paul R Milgrom, "The limited influence of unemployment on the wage bargain," *American economic review*, 2008, *98* (4), 1653–74.

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

Harrington, Joseph E. and Andrzej Skrzypacz, "Private Monitoring and Communication in Cartels: Explaining Recent Collusive Practices," *The American Economic Review*, 2011, *101* (6), 2425–2449.

Harstad, Bård, "Climate contracts: A game of emissions, investments, negotiations, and renegotiations," *Review of Economic Studies*, 2012, *79* (4), 1527–1557.

___ , "The dynamics of climate agreements," *Journal of the European Economic Association*, 2016, *14* (3), 719–752.

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

___ **and** ___ , "Property Rights and the Nature of the Firm," *Journal of political economy*, 1990, *98* (6), 1119–1158.

Jervis, Robert, "Cooperation under the security dilemma," *World politics*, 1978, *30* (2), 167–214.

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

Klimenko, Mikhail, Garey Ramey, and Joel Watson, "Recurrent trade agreements and the value of external enforcement," *Journal of International Economics*, March 2008, *74* (2), 475–499.

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.

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.

Ramey, Garey and Joel Watson, "Contractual intermediaries," *Journal of Law, Economics, and organization*, 2002, *18* (2), 362–384.

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

Safronov, Mikhail and Bruno Strulovici, "Contestable Norms," 2018.

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

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.

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

van der Ploeg, Frederick and Aart J de Zeeuw, "Perfect equilibrium in a model of competitive arms accumulation," *International Economic Review*, 1990, pp. 131–146.

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.

____ , "Contract and game theory: Basic concepts for settings with finite horizons," *Games*, 2013, *4* (3), 457–496.

____ **, David A. Miller, and Trond E. Olsen**, "Relational Contracting, Negotiation, and External Enforcement," *American Economic Review*, 2020, *110* (7), 2153–97.

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