Soft-Landing Contracts and Temporal Incompleteness*

Bernhard Ganglmair† Julian Klix‡ Dongsoo Shin§

March 15, 2024

Preliminary version!

Abstract

In two-player finite-horizon repeated games, end-of-game behavior usually leads to patterns of decreasing degrees of cooperation with diminishing payoffs. We study (incomplete) contracts that restrict the action space only in later stages of the players' relationship. These soft-landing contracts complement the relational contract dynamics in the early stages of the relationship and are as effective in inducing cooperation as contracts with full commitment while being more cost-effective. We find soft-landing contracts to be able to increase cooperation both in the intensive and extensive margin, with our results proving robust to extensions of the model regarding more complex cost structures and endogenous timing of contracting.

Keywords: cooperation, incomplete contracts, relational contracts, repeated games

JEL Codes: C73, D86, L14

*We thank Nicolas Schutz for helpful comments and suggestions. Support by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through CRC TR 224 (Project B02) is gratefully acknowledged. The authors declare that they have no relevant or material financial interests that relate to the research described in this paper.

†University of Mannheim and ZEW Mannheim, Germany. E-mail: b.ganglmair@gmail.com

‡University of Mannheim, Germany. E-mail: julian.j.klix@gmail.com

§Santa Clara University, Santa Clara, USA. E-mail: dshin@scu.edu

1 Introduction

Contracts are seen as having a central role in generating cooperation in business and other relationships. Although it is one of the most widely used commitment devices in various transactions, contracting is not free. It requires resources—from a transaction cost perspective, contracts are associated with costs of designing and monitoring. As such, contracting alone is insufficient to accommodate many important economic phenomena. The role of trust in contractual relationships, in particular, has been emphasized in various studies (e.g., Burchell and Wilkinson, 1997). Practitioners frequently point out that being better at using contracts starts with good strategies to foster trust (e.g., Tippins, 2021). That is, such strategies do more than just reduce the cost of contracting—building trust and using contracts are complementary to each other.

As reported by Ohmae (1989) in the Harvard Business Review, all business agreements may not last long, but many business alliances depend on "softer" aspect of relationships. He notes that few alliances, including those engaged in joint ventures and supply chains of various industries, rely on rigidly binding agreements from the beginning, stressing the importance of non-binding agreements on two-sidedness or symmetry of effort and attention in earlier stages to solidify their relationships in later stages.

In this paper, we study how a non-binding agreement in an earlier stage of the game can be used to induce cooperative behavior in a baseline with, at most, partial cooperation. This temporally incomplete contract is as effective as a more standard complete contract while being favorable when it comes to contracting costs. We show that such *soft-landing contracts* are able to increase cooperation in both the intensive and extensive margin while being optimal compared to both the complete contract and a baseline without contracts.

We start with a simple, symmetric, finitely repeated game that can generate partial cooperation in the early stages. To that end, our stage game, based on a prisoner's dilemma game, features three actions – keeping it as concise as possible while avoiding complete backward unraveling (and thus complete lack of cooperation) that would be inherent to two-action stage games. A *grim trigger* strategy that unforgivingly punishes deviation then generates a cooperative equilibrium where players cooperate in all early stages as long as they do not discount the future too much.

We then introduce contracts in the form of commitment devices that allow players to commit to playing (or avoiding) specific actions with the goal of extending the stages in which cooperation takes place (*intensive margin*) or the parameters in which cooperation can occur entirely (*extensive margin*). We find that while simply prescribing the optimal action in all periods (i.e., a full commitment contract) is effective in achieving both goals, once costs of contracting are factored in, less prescriptive contracts are superior in many circumstances, in particular when cooperation could already be achieved partially without contracts.

These *soft-landing contracts* are characterized by their temporal incompleteness; leaving the players free hand in the early stages and only coming into effect in the late stages of the game. This way, they are able to prevent the backward unraveling that precludes late-stage cooperation without eliminating the incentives used by the grim trigger strategies that generate cooperation already in the baseline. While also requiring a maximum discounting, this is less restrictive than the baseline requirements, leading to situations in which the contracts can induce cooperation extensively, i.e., when there has not been cooperation before, as well as intensively —increasing the periods with cooperation. Due to its simpler and shorter nature, these soft-landing contracts are less costly than the complete full commitment alternative. Hence, we find that whenever they are profitable (i.e., better than the contract-less baseline), they are also optimal. These results are robust to richer models that incorporate more complex cost structures, endogenize the negotiation timing, or relax the assumptions of prescriptive contracts in favor of contracts over monetary transfers.

Our analysis contributes to the literature on incomplete contracts, particularly on the question of when and why parties choose to write more or less complete contracts. Contractual incompleteness is often associated with transaction costs (Williamson, 1985, 1989) that can arise both ex ante and ex post. Ex-ante transaction costs are the costs associated with the drafting of the contract ("search costs" (e.g., Klein, 2002; Tirole, 2009) and "ink costs" (Dye, 1985; Anderlini and Felli, 1994; Melumad et al., 1997; Battigalli and Maggi, 2002, 2008)). Ex-post transaction costs are the costs of enforcing and implementing contracts. Kaplow (1995) and Schwartz and Watson (2004) argue that more complex contracts (or rules) may be more costly to enforce (e.g., litigate) because more information or evidence is needed. Furthermore, more detailed contracts specifying more contingencies imply higher costs of monitoring compliance and detecting violations—as highlighted by Diamond (1984, 1991) in the context of corporate debt.[1] In our model, we assume fixed ex-post enforcement costs (but no ex-ante drafting costs) that do not depend on the characteristics or duration of the contract. Contracting parties incur these costs only once, and because of time discounting, they prefer contractual commitments in later periods. This adds a temporal dimension to optimal contractual incompleteness; enforcement costs of longer and earlier contracts are effectively higher.

More detail is only one aspect of completeness. Battigalli and Maggi (2002) argue that there may be tradeoffs between additional detail and overly rigid contracts where contractual contingencies may be insufficiently customized to the contracting parties. They show that when the surplus from a contractual relationship relative to the contracting costs is sufficiently small, the optimal contract (signed by the parties) leaves some degree

[1]Also, in multi-tasking environments, imperfect monitoring can result in the optimal contract being simple (Holmstrom and Milgrom, 1991).

of discretion to the party—that means, the contract is incomplete as some state-action pairs are not specified. Basov (2016, 62–64) provides a simple example to illustrate this result. In line with these results, we show that flexible soft-landing contracts with a larger allowable action space outperform rigid soft-landing contracts that effectively prescribe the cooperative action.

Our baseline setup with cooperation in a finitely repeated game is related to work in the literature on relational contracts. In finitely repeated games, cooperation is typically not attainable because of the underlying end-of-game properties. Yet, experimental evidence shows that cooperation is possible in such setting (Axelrod, 1981). Kreps et al. (1982) establish cooperation in a finitely repeated prisoner's dilemma by introducing reputation and asymmetric information. We take the approach in Benoit and Krishna (1985) to achieve some cooperation in our baseline model. We contribute to the literature by combining relational and formal contracting. We show that formal contracts can complement relational contracts.

In addition, our paper relates to the literature on contracting practice where it has long been established that negotiating (and upholding) contracts is intricately linked to questions of trust. Arrighetti et al. (1997) observe how contracting and trust interact in the environment of contracting law, and a wide array of papers has studied how contracts can be used to form trust (e.g., Ben-Ner and Putterman, 2009; Malhotra and Murnighan, 2002; Burchell and Wilkinson, 1997). However, trust can also play an instrumental role in shaping contracts. In an experimental setting, Ben-Ner and Putterman (2009) find that complete contracts tend to substitute trust, while—related to the soft-landing contracts we are studying—Lorenz (1999) concludes that trust plays an important role in forming incomplete contracts. Similarly, Burchell and Wilkinson (1997) observe in a survey on European firms a positive relationship between trust and flexibility outside the contract. More specifically, Ohmae (1989) and Williamson (1985, 1989) note how relational aspects such as trust lead to contracts that often start as non-binding agreements or unspecific framework agreements before becoming more binding and specific in the course of the interaction.

Our model does not feature trust (but rational, profit-maximizing players instead) and thus provides an alternative explanation for the soft-landing contracts observed in practice. We show how this kind of temporally incomplete contract can also be viewed as a strategic choice in addition to merely an expression of trust between contracting parties.

The rest of the paper is structured as follows: In Section 2, we analyze the baseline model and characterize the benchmark equilibria without contracts. In Section 3, we introduce soft-landing contracts and derive the parameter space in which they can increase cooperation relative to the benchmark equilibria. In Section 4, we characterize the most preferable contractint solution that maximizes joint profits. In Section 5, we extend the

		Player P2		
		w (work)	l (loaf)	s (shirk)
Player P1	w (work)	$\begin{pmatrix} 0 \\ 0 \end{pmatrix}$	$\begin{pmatrix} \tilde{e}\pi - m \\ \tilde{e}(\pi - c) \end{pmatrix}$	$\begin{pmatrix} 2\pi - m \\ 2(\pi - c) \end{pmatrix}$
	l (loaf)	$\begin{pmatrix} \tilde{e}(\pi - c) \\ \tilde{e}\pi - m \end{pmatrix}$	$\begin{pmatrix} \tilde{e}/2 \\ \tilde{e}/2 \end{pmatrix}$	$\begin{pmatrix} (2+\tilde{e})\pi - c - m \\ (2+\tilde{e})\pi - 2c \end{pmatrix}$
	s (shirk)	$\begin{pmatrix} 2(\pi - c) \\ 2\pi - m \end{pmatrix}$	$\begin{pmatrix} (2+\tilde{e})\pi - 2c \\ (2+\tilde{e})\pi - c - m \end{pmatrix}$	$\begin{pmatrix} 1 \\ 1 \end{pmatrix}$

Table 1: Normal-Form Representation of the Stage Game

baseline model to allow for more flexible timing and contracting costs. In Section 6, we conclude.

2 Baseline Model of Interaction

2.1 Setup

Let there be two players interacting in a finitely repeated normal-form game. In order to obtain early cooperation, we focus on a repeated game in which there are three possible actions, yielding an *augmented prisoner's dilemma* (APD) stage game. This approach allows using a grim trigger strategy of Nash reversion to sustain the cooperative outcome in early periods.[2]

The stage game consists of two players (P1 and P2) deciding whether to work (w), loaf (l), or shirk (s),[3] associated with effort levels $e_i \in \{2, \tilde{e}, 0\}$ corresponding to $A_i = \{w, l, s\}$. The payoff in each stage game outcome is then given by

$$\pi_i(a_i, a_j) = (a_i + a_j)\,\pi - a_i c - \mathcal{I}(a_i < a_j) \cdot m$$

where π is the utility of joint effort, c is the cost of effort, and m is the moral utility cost of free-riding. To simplify the payoff structure, we normalize the payoffs for the (w, w) outcome to 1 by setting π and c such that $\pi_i(w, w) = 4\pi - 2c = 1$. We summarize the symmetric payoff matrix of the stage game in Table 1.

To ensure that (s, s) and (l, l) form the only two pure-strategy Nash equilibria in the

[2]This follows the limited folk theorem using trigger strategies proposed by Benoit and Krishna (1985) (see also Fudenberg and Tirole (1991)).

[3]In this game w dominates s, whereas the intermediate action l is neither dominant nor dominated. Thus, we refer to this structure as an *augmented prisoner's dilemma*

stage game, we impose the following conditions.[4]

Assumption 1. *Let the set of parameters* $(\pi, c, m, \tilde{e})$ *be such that (a)* $\pi > \frac{1}{2}$, *(b)* $c = 2\pi - \frac{1}{2}$, *(c)* $0 < m < 2\pi - 1$, *(d)* $0 < \tilde{e} < \min\left(\frac{m}{\pi - 1/2}, \frac{m+1/2}{\pi}\right)$.

For better tractability, let $\pi^D = \pi_i(s, w) = 2\pi - m$ denote the maximal deviation payoffs and $\pi^L = \pi_i(l, l) = \tilde{e}/2$ denote the payoffs of the (l, l)-NE.

The APD stage game is repeated T times with a standard discounted total utility $\Pi_i = \sum \delta^t \pi_i(a_{i,t}, a_{j,t})$. We make the standard assumptions of rational players and common knowledge of past actions.

2.2 Baseline Equilibria

Because the APD stage game has two pure-strategy Nash equilibria, consider the following strategy to form a *cooperative equilibrium*: If everyone cooperates, cooperation continues. If someone deviates, (s, s) will be played as a punishment for the remainder of the game. Formally, this is saying every player plays the grim trigger strategy

$$a_{i,t}^{GT} = \begin{cases} s & \text{if } \exists \tau < t : a_{j,\tau} \neq w; \\ l & \text{if } t = T \wedge \forall \tau < t : a_{j,\tau} = w; \\ w & \text{otherwise.} \end{cases}$$

Note that because payouts in each stage game are independent, by backward induction, an equilibrium has to be played in the last period. The outcome in the last period is, therefore, (l, l) rather than (w, w) even if both players cooperated up until the last period. The timing of the optimal deviation then depends on the discount factor δ, with a last-period deviation being optimal whenever $\delta > \frac{\pi^D - 1}{\pi^D}$.[5] Factoring in discounting, a deviation is unprofitable whenever

$$\frac{1 - \delta^{T-1}}{1 - \delta} + \delta^{T-1} + \delta^T \pi^L > \frac{1 - \delta^{T-1}}{1 - \delta} + \delta^{T-1} \pi^D \iff 1 + \delta \pi^L > \pi^D \iff \delta > \frac{\pi^D - 1}{\pi^L}$$

This relationship yields the first equilibrium condition, summarized in the following lemma.

Lemma 1 (Cooperative Equilibrium)**.** *A cooperative equilibrium in which* $a_{i,t}^{GT}$ *is played yielding* $(w, w), \ldots, (w, w), (l, l)$ *exists in the baseline model whenever* $\delta > \dfrac{\pi^D - 1}{\pi^L} =: \bar{\delta}^{BL}$.

[4] These assumptions also guarantee the existence of only one mixed-strategy Nash equilibrium in the stage game. For better tractability, we focus our analysis solely on pure-strategy Nash equilibria. Allowing for mixed strategies in our analysis, however, does not change the results qualitatively. For a discussion of the effects of mixed-strategy equilibria, see Appendix A.1. Alternatively, the existence does not affect the outcomes under the modified assumption $(d)'$ $\frac{m}{\pi} < \tilde{e} < \min(\frac{m}{\pi - 1/2}, \frac{m+1/2}{\pi})$.

[5] The derivation of the thresholds for the APD can be found in Appendix A.1.

In addition to the cooperative equilibrium, it is also always optimal (in the sense of a *subgame-perfect Nash equilibrium*) to play a Nash equilibrium in every period. That means that every action profile $a_1^* = a_2^* \in \{s, l\}^T$ also constitutes a subgame-perfect Nash equilibrium of the APD, of which $a_i^* = (l, \dots, l)$ is payoff-dominant, forming the *non-cooperative equilibrium* characterized in Lemma 2.

Lemma 2 (Non-Cooperative Equilibrium)**.** *A non-cooperative equilibrium in which $a_{i,t} = l$ is played yielding $(l, l), \dots, (l, l)$ exists in the baseline model for any δ.*

This equilibrium is of particular interest in the case when cooperative equilibria cannot be sustained, that is, whenever $\delta < \bar{\delta}^{BL}$.

3 Soft-Landing Contracts

While limited cooperation can be sustained in the baseline model, adding the ability to write contracts to the framework will give the players tools to improve cooperation—both *intensively* (increasing the number of periods with cooperation) and *extensively* (extending the region of the parameter space in which cooperation is possible). Assuming that actions are verifiable (and contractible), we model a contract in this context as a commitment device that allows the players to commit to a certain subset of actions available to play for a specified subset of periods.

The allowable actions from this contract can, but do not have to, depend on earlier actions taken. For our main results, we assume that this commitment is made (and the contract entered) before the first-period interaction. At the outset of the game ($t = 0$), players commit to a contract by specifying a subset of actions available to them at a later stage, that is, when the first contracted period is reached.[6] Formally, a contract can be specified as $C = (C_1, C_2)$ with $C_1 = C_2$ where $\tilde{A}_i(h) = C_i(h) \in (\mathcal{P}(A_i)/\emptyset)$ is the set of actions available to player i at history h. Later in the paper, we relax this pre-commitment assumption and allow for endogenous contract timing (in $t \geq 1$) after the interaction has begun.

Contracting comes at a cost. We assume that these costs are *implementation costs* $\kappa \geq 0$ that are incurred once at the first contracted period $t_\kappa := \min\{t : C_t \neq A_t\}$ (rather than at the contracting stage).[7] We assume the costs κ are independent of the length

[6]This is the same as saying that the contract entered in $t = 0$ does not restrict the action space nor does it prescribe any actions between the contract stage and the first contracted period.

[7]Implementation costs are associated with the costs of enforcing a contract (e.g., Kaplow, 1995; Schwartz and Watson, 2004) or the costs of monitoring compliance and detecting violations (e.g., Diamond, 1984, 1991). In an extension (see Section 5.2), we will consider more general contracting costs that also include drafting costs (incurred at the time the contract is entered), associated with the costs that arise from researching and analyzing contingencies (*search costs*) (e.g., Klein, 2002; Tirole, 2009) and from specifying these contingencies (*ink costs*) (e.g., Dye, 1985; Anderlini and Felli, 1994; Melumad et al., 1997).

of the contract, the number of contractual periods, or the amount of committed actions. These assumptions make our model more tractable and eliminate potential biases in favor of shorter and simpler soft-landing contracts.

In this section, we establish the *implementability* of cooperation with soft-landing contracts for three scenarios: (a) full commitment, in which players commit to actions in every period; (b) history-independent soft-landing contracts, in which committed actions are not contingent on the history of actions prior to the contract; and (c) history-dependent contracts, in which the contracted actions depend on the history leading up to the contract. We return to the question of profitability (when the contract strictly payoff-dominates the relevant baseline scenario from the previous section) and optimality (when the contract also strictly payoff-dominates all other contracts) of a soft-landing contract in Section 4.

3.1 Full Commitment

The full commitment contract scenario represents the most restrictive commitment and acts as a benchmark; it is the polar opposite of the environment without any commitment in our baseline equilibria. In this full commitment scenario, both players commit to a restricted action space $\tilde{A}_i$ for every period. Trivially, a full commitment contract that restricts the action space to $\tilde{A}_i = \{w\}$ and effectively prescribes $a_{i,t} = w$ for all i and all t implements full cooperation for the entirety of the extended game. Trivially, this cooperation equilibrium under a full commitment contract exists regardless of the specific δ. This sets this contract apart from the contracts we consider next.

3.2 History-Independent Soft-Landing Contracts

In a soft-landing contract, the players commit to a (restricted) action space $\tilde{A}_i$ (effectively prohibiting some actions) for the latter part of their interaction.[8] With history independence, the action space $\tilde{A}_i$ does not depend on past actions. For simplicity, we require the restrictions on the action space to be symmetric.

We now construct a soft-landing contract and players' strategy under this contract that can implement the fully cooperative outcome. For simplicity, we consider a contract that does not interfere for the first $T-1$ period and only takes effect in the last period. It prohibits the lowest-effort action s by specifying a (restricted) action space $\tilde{A}_i = \{l, w\}$.[9]

[8] We assume players abide by the rules of the contract, which means do not choose the eliminated action. This can achieved by including appropriate penalties. We do not specify such contracted penalties but simply assume that a violation of the contract terms is prohibitively costly.

[9] The arguably more natural case of an *abridged full commitment contract* that prescribes $a_i = w$ (with action space $\tilde{A}_i = \{w\}$) only for the last τ periods (which is close in spirit to the full commitment contract), is strictly dominated by the proposed soft-landing contract. It effectively partitions the game into forced cooperation for the last τ periods with a shortened ADP repeated $T-\tau$ times. Analogous to the discussion for the baseline in Lemma 1, cooperation cannot be sustained in period $T-\tau$, which

With this contract, the (restricted) stage game becomes

	loaf (l)	work (w)
loaf (l)	(π^L, π^L)	(π_{wl}, π_{lw})
work (w)	(π_{lw}, π_{wl})	$(1, 1)$

This stage game continues to have two pure-strategy Nash equilibria, namely (l, l) and (w, w). This is by Assumption 1 that guarantees that $\pi_{wl} < 1$ and $\pi_{lw} < \pi^L$.

Under the soft-landing contract and the (restricted) stage game in the last period, let the players' grim-trigger strategy be as follows. It is analogous to the cooperative strategy in the baseline scenario but, in the last period, rewards cooperation with (w, w) and punishes deviation with (l, l). This strategy is effectively shifting both punishment and reward upwards payoff-wise. Formally, we define the strategy as:

$$a_{i,t}^{IA} = \begin{cases} s & \text{if } t \neq T \land \exists \tau < t : a_{j,\tau} \neq w \\ l & \text{if } t = T \land \exists \tau < t : a_{j,\tau} \neq w \\ w & \text{otherwise} \end{cases}$$

We have an optimality of cooperation with this new grim trigger strategy whenever

$$\frac{1-\delta^{T-1}}{1-\delta} + \delta^{T-1} + \delta^T > \frac{1-\delta^{T-1}}{1-\delta} + \delta^{T-1}\pi^D + \delta^T \pi^L \iff \delta > \frac{\pi^D - 1}{1 - \pi^L} =: \bar{\delta}^{IA}.$$

We summarize this result of cooperation under the soft-landing contract in the following lemma.

Lemma 3 (History-Independent Soft-Landing Contract)**.** *A cooperative equilibrium in which $a_{i,t}^{IA}$ is played yielding $(w, w), \ldots, (w, w)$ exists whenever $\delta > \frac{\pi^D-1}{1-\pi^L} =: \bar{\delta}^{IA}$ if a soft-landing contract prescribing history-independent actions is used.*

Crucially, this contract works—in contrast to an abridged full commitment— because while eliminating the most profitable deviation action s it is not too rigid to eliminate the flexibility of punishing previous behavior. Put differently, SLCs that are not themselves contingent, only work if not too rigid, see also Battigalli and Maggi (2008).

3.3 History-Dependent Soft-Landing Contracts

Next we turn to contracts that can be conditioned on the history of past play, i.e., their effect depends on what actions players took in the past. Even though the history-independent contracts discussed before can be understood technically as degenerate history-dependent contracts, this section will focus on proper history-dependent contracts.

is discounted less and thus leads to a lower present value utility (even at $\kappa = 0$).

These contracts can be used to implement a strategy that employs otherwise incredible punishments because the requirement for being best responses disappears while the necessary requirements for punishments to be conditioned on past behavior is now possible.
For comparability and tractability consider a contract that, similar to the history-independent case, only lasts for the last period.[10] To deter deviation, let it be such that it prescribes w as long as the opponent played w in all past periods and s otherwise.[11] This results in the modified grim trigger strategy $a_{i,t}^{DA}$ being played:

$$a_{i,t}^{DA} = \begin{cases} s & \text{if } \exists \tau < t : a_{j,\tau} \neq w \\ w & \text{otherwise} \end{cases}$$

Note that except for the last period the standard grim trigger strategy $a_{i,t}^{GT}$, already implements the most effective (i.e. the harshest) punishment.
Then with deviation taking place in period $T-1$, deviation would be unprofitable whenever

$$\frac{1-\delta^{T-1}}{1-\delta} + \delta^{T-1} + \delta^{T} > \frac{1-\delta^{T-1}}{1-\delta} + \delta^{T-1}\pi^D + \delta^T 0 \iff \delta > \pi^D - 1 =: \bar{\delta}^{DA}$$

as summarized by Lemma 4 below

Lemma 4. *A cooperative equilibrium exists whenever $\delta > \pi^D - 1 =: \bar{\delta}^{DA}$ if a SLC prescribing history-dependent actions is used.*

4 Results

In this section we compare the contracts described in the previous section to the baseline (profitability) as well as to each other (optimality). To this end, we will assume that the payoff-dominant equilibrium is played in the baseline.[12] Again, by focusing on the payoff dominant equilibrium in the baseline we eliminate biases in favor of contracts when comparing to the baseline.

4.1 Existence of Cooperative Equilibria

As we briefly discussed in Section 2 and 3 all equilibria have conditions on δ for their existence. To better illustrate this conditions for the existence of cooperative equilibria in the different frameworks, we consider a specific subset of the parameter space created

[10] A more general analysis can be found in Section 5.1

[11] Formally, this is saying $C_i(h) = \{w\}$ if $h = (w,w)^{T-1}$ and $C_i(h) = \{s\}$ otherwise.

[12] While in general the existence of a payoff ordering of equilibria does not ensure the payoff dominant equilibria to be played, one can think of the contracting stage additionally as a coordination instrument for the baseline.

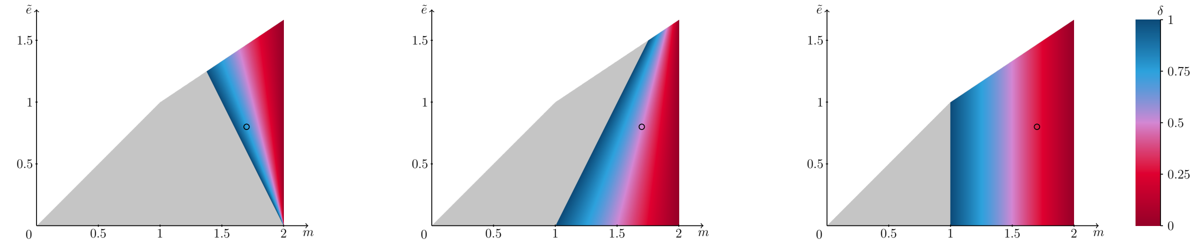

(a) Baseline (b) History-Independent (c) History-Dependent

Figure 1: Heatmaps for the cooperative equilibrium threshold for $\bar{\delta}^{BL}$ (a), $\bar{\delta}^{IA}$ (b) and $\bar{\delta}^{DA}$ (c) for $\pi = 1.5$.

by Assumption 1. Figure 1 illustrates the thresholds for the baseline model as well as the history-independent and dependent SLC as summarized in Lemma 1, 3 and 4 respectively.

While it is clearly visible that each cooperative equilibrium is feasible for a range of parameters of the parameter space, it is equally discernible that there is no complete dominance ordering of environments. In other words, there are different parameters for which different types of environments can produce a cooperative equilibrium. This comparative result is emphasized by Figure 2.

There is however partial ordering regarding the history-dependent SLCs. Recall that the history-independent SLC is imposed by merely shifting the action profiles used for reward and punishment upwards in terms of payoff. In contrast, history-dependent SLCs optimize the grim-trigger strategy used in the baseline. As a consequence, it is possible to show that whenever a cooperative equilibrium is feasible in the baseline, so is cooperation under a history-dependent SLC.[13] The same is not true for between the baseline and history-independent SLCs as the exemplary parameter-subspace depicted in Figure 2 shows.

Further, compare the history-dependent and independent SLCs by looking at Lemmas 3 and 4. It becomes apparent that whenever a history-independent contract can induce a cooperative equilibrium, so can a history-dependent SLC.[14] This should not come as a surprise given that history-independent SLCs form a subcategory of the more general set of history-dependent SLCs. This result is summarized in Corollary 1.

Corollary 1. *The feasibility of a cooperative equilibrium under a history-independent SLC implies the existence of a cooperative equilibrium with a history-dependent SLC, which is equivalent to* $\bar{\delta}^{DA} < \bar{\delta}^{IA}$.

While all equilibria looked at so far require a certain level of discounting, the full commitment contract described in Section 3.1 is a special case. Since players commit to one action for the entire game with no opportunity to deviate, full commitment contracts are intrinsically feasible for any δ.

[13]This follows simply from $\pi^L < 1$, which implies $\bar{\delta}^{BL} := \frac{\pi^D - 1}{\pi^L} > \pi^D - 1 =: \bar{\delta}^{DA}$.

[14]This follows directly from the fact that $\bar{\delta}^{DA} = \pi^D < \pi^D / 1 - \pi^L = \bar{\delta}^{IA}$ since $\pi^L \in (0, 1)$

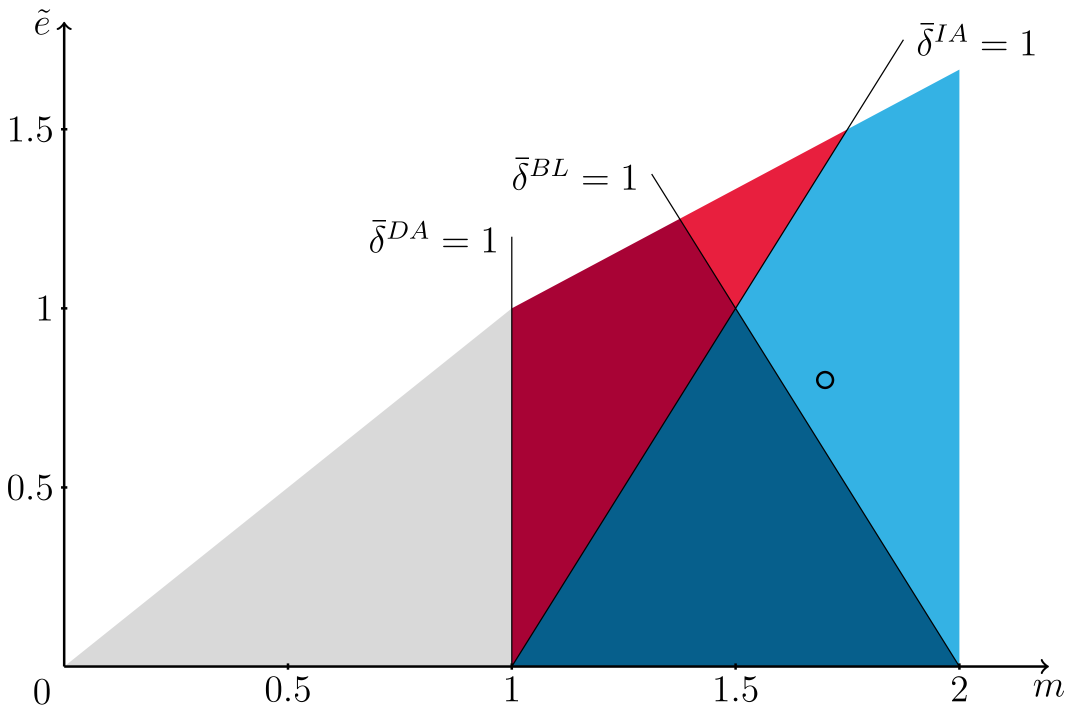

Figure 2: Overview over the different types of implementable contracts for $\pi = 1.5$. The corresponding value of the discount threshold $\bar{\delta}^i$ is represented by the heatmap (with no implementability in the grey area). Example 1 is denoted by the circle.

4.2 Profitability of Contracts

This dominance in implementability of the full commitment contract is, however, only one half of the story. To get the full picture, we need to turn to the costs and the implementability of the SLCs with regards to payoffs. This question has to be examined in comparison to the payoff of the baseline equilibria.

Recall that the cooperative equilibrium of the baseline model is characterized by a payoff stream of $\pi_i(w, w) = 1$ in the first $T - 1$ periods and $\pi_i(l, l) = \pi^L$ in the last. The non-cooperative equilibrium on the other hand pays out $\pi_i(l, l) = \pi^L$ in every period. Recall further that contractual costs were modelled as implementation costs, i.e. the cost κ is paid in the first period the contract comes into effect.

Now consider the full commitment contract described before. It yields a payoff of $\pi_i(w, w) = 1$ in every period and is, thus, *ex-ante* preferable to the cooperative equilibrium whenever

$$\frac{1-\delta^T}{1-\delta} + \delta^T - \kappa > \frac{1-\delta^T}{1-\delta} + \delta^T \pi^L \iff \kappa < \delta^T(1-\pi^L)$$

Intuitionally, since the full contract only changes the payoff of the last period the cost of the contract can only be as high as the discounted difference in payoffs. In contrast, to be preferable to the non-cooperative equilibrium it has to be that

$$\frac{1-\delta^T}{1-\delta} + \delta^T - \kappa > \left(\frac{1-\delta^T}{1-\delta} + \delta^T\right)\pi^L \iff \kappa < \left(\frac{1-\delta^{T+1}}{1-\delta}\right)(1-\pi^L)$$

Note that this directly implies a higher cost threshold in the case of a no-cooperation baseline. Thus there are contract costs κ for which full commitment contracts are only profitable if there was no cooperation before. These results are summarized in the following Proposition:

Proposition 1. *As long as* $\kappa < \frac{1-\delta^{n+1}}{1-\delta}(1-\pi^L)$ *the cooperative equilibrium induced by a full commitment contract is preferable to the non-cooperative baseline. Whenever* $\kappa < \delta^T(1-\pi^L)$ *it is also preferable to the cooperative baseline.*

Next we turn to the history-independent SLC that prohibits the strongest deviation. As before this results in a payoff stream of $\pi_i(w,w) = 1$ in every period. The major difference, however, is that this SLC only comes into effect in the last period of the game. Hence, the implementation of such a class of SLC is optimal compared to the (cooperative) baseline whenever

$$\frac{1-\delta^T}{1-\delta} + \delta^T(1-\kappa) > \frac{1-\delta^T}{1-\delta} + \delta^T\pi^L \iff \kappa < 1-\pi^L$$

since the costs κ are paid in the period in which the contract first becomes binding, the last in this case. As in the full contract case, a SLC is also possible in a case with a non-cooperative equilibrium in the baseline. Then the cost-threshold becomes $\kappa < (1+\ldots+\delta^{-n})(1-\pi^L)$, which is clearly higher compared to the case of a cooperative equilibrium in the baseline. Proposition 2 summarizes those two cases:

Proposition 2. *As long as* $\kappa < \frac{1-\delta^{n+1}}{(1-\delta)\delta^T}(1-\pi^L)$ *the cooperative equilibrium induced by a history-independent SLC is preferable to the non-cooperative baseline. Whenever* $\kappa < 1-\pi^L$ *it is also preferable to the cooperative baseline.*

This has two implications: Firstly, as long as $\kappa < \delta^T(1-\pi^L)$, both a full contract and a SLC are possible, but the SLC is strictly preferred to the full contract. Secondly, if $\delta^T(1-\pi^L) < \kappa < 1-\pi^L$, only the SLC but not the full contract is profitable.

Lastly, there is the class of history-dependent SLCs. Since the payoff stream induced by the cooperative equilibria is identical to the case of history-independence, the cost conditions remains $\kappa < 1-\pi^L$. The same is true for SLCs in a non-cooperative baseline setting with $\kappa < (1+\ldots+\delta^{-n})(1-\pi^L)$ as the profitability threshold. These two thresholds are reflected in the following proposition

Proposition 3. *As long as* $\kappa < \frac{1-\delta^{n+1}}{(1-\delta)\delta^T}(1-\pi^L)$ *the cooperative equilibrium induced by a history-dependent SLC is preferable to the non-cooperative baseline. Whenever* $\kappa < 1-\pi^L$ *it is also preferable to the cooperative baseline.*

4.3 Optimality of Contracts

Having considered the profitability of the different contracts, the question that comes next naturally is which contract is optimal from the players' *ex-ante* perspective. For that we compare the net payoffs between the different types of contracts. We will compare the three main contracts discussed so far; a contract featuring full commmitment (FC), one prescribing independent actions (IA) and the one prescribing dependent actions (DA). Starting with a direct comparison between IA and DA, note that since the cost structure of both SLC types is identical, all results for the history-independent contract with regards to cost translate to the history-dependent contracts, reflective in the equivalence of Propositions 2 and 3. Since the cost-thresholds are identical, resulting in the same implication for both contract types if feasible. A further comparison with the full commitment contract (1) shows that, since costs have to be paid already in the first rather than the last period, any of the two contracts is preferable if available.

This result goes to show a general intuition as to why SLCs are a reasonable contract practice that might be preferable not only because of shorter contractual duration and complexity, but also from a purely payoff perspective. Games (or more generally interactions) between players that are approached by backwards induction often feature *unravelling*. As a result, cooperation (if feasible at all) takes place in the first periods before purely self-centered behavior takes precedence in the last stages.

With the goal of ensuring cooperation, the first (in this model $T-1$) periods already achieve that without the interference of any contract. Thus, the most duration-efficient contract would be one that only intervenes in periods without prior cooperation in order to induce additional cooperative periods. This aspect of a SLC increasing the amount of already existing cooperation periods will be referred to as inducing cooperation on the *intensive margin*. Since costs are incurred when the contract first takes effect and players discount the future, duration efficiency (with respect to early periods) translates directly into payoff efficiency. For that reason any feasible SLC necessarily payoff-dominates the full commitment contract.

However, Section 3.2 also showed that such a SLC can not just simply be an abridged version of the full commitment contract as it would just move the time-point of deviation forward without any punishment in the contractual periods. In contrast, we presented an alternative for constructing a contract that is still independent of previous actions but retains the possibility of punishing players while the contract is active. To understand the intuition behind that approach, it is necessary to understand that by restricting the contract to be history-independent, it is deprived of the means to inflict punishment for deviating prior to the contract becoming active. Thus, punishment needs to take place in the subset of actions the contract still allows the players to play. In the case of committing to one single actions for a period, cooperation and deviation in any prior period can only be met by the same prescribed action. If the contract leaves the option of more than one action, however, players retain the possibility to react differently to cooperation or deviation by adhering to, for instance, a modified grim-trigger strategy. Thus, the SLC is constructed by merely prohibiting the action resulting in the most profitable deviation in the only period in which no cooperation takes place in the baseline. This is clearly not necessary to take into account when constructing a SLC that can be conditioned on prior play, which thus allows to implement an optimal grim-trigger strategy to induce cooperation.

There is another dimension of inducing cooperation not yet touched upon, the *extensive margin*. In addition to increasing the periods of already existing cooperation, a contract can also induce cooperation by changing the equilibrium from a non-cooperative one in the baseline to a cooperative one. All SLCs are able to extensify cooperation in that manner. However, they are limited by their implementability constraint on the discount factor, $\bar{\delta}^{IA}$ and $\bar{\delta}^{DA}$ respectively. The same is not true for the full commitment contract. By its nature it prevents deviation entirely and is thus not dependent on the discount factor. Given low enough contractual costs κ, while suboptimal for intensifying cooperation from a payoff perspective, a full commitment contract is always able to extend cooperation when the SLCs are not. While this occasional trade-off between the extensive and intensive margin in the form of choosing either a SLC or a full commitment contract is apparently rather extreme, Section 5.1 explores the middle ground of multi-period SLCs creating a more gradual trade-off between the payoff-efficiency of SLCs and the feasibility advantages of long contracts.

To illustrate these prior points consider one specific example:

Example 1. *Let the stage game described in Section 2 be repeated $T = 3$ times with $\pi = 1.5, m = 1.7, \tilde{e} = 0.8$.*

As evident in Figure 2, this example is such that both types of SLC, history-dependent and independent, can be implemented. Figure 3 analyzes this specific example further, focusing not only on implementability but also on payoff-optimality by including the contractual costs κ in the parameter space. In this specification whenever $\delta > \bar{\delta}^{BL} = 0.75$ (marked by the vertical black line) a cooperative equilibrium can be sustained in the baseline. If additionally $\kappa < 1 - \pi^L = 0.6$ both SLC types yield a higher payoff than the baseline. In this case, since there was already a cooperative equilibrium in the first place, cooperation is induced in the *intensive margin*. Consider in contrast the areas to the left of the black line. Here $\delta < \bar{\delta}^{BL}$ and thus the baseline can only feature a non-cooperative equilibrium. Consequently, any contract that induces cooperation does so in the *extensive margin*. In this parametrization $\bar{\delta}^{IA} = 0.5 > 0.3 = \bar{\delta}^{DA}$ so that there exists a low range of discount factors for which only a history-dependent SLC can induce cooperation.[15] Costs

[15]While we showed before that necessarily $\bar{\delta}^{DA} < \bar{\delta}^{BL}$ depending on the parametrization $\bar{\delta}^{IA} \gtrless \bar{\delta}^{BL}$,

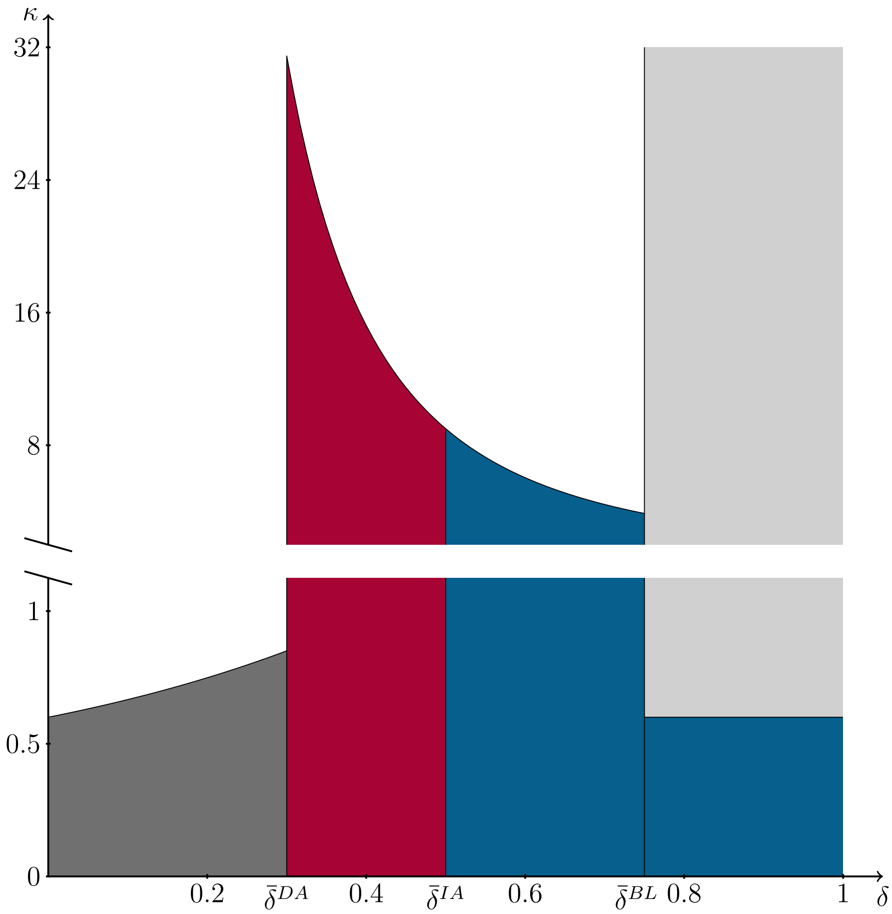

Figure 3: Environments with the highest payoff depending on discount δ and cost κ for the specific example marked before (with $\pi = 1.5, m = 1.7, \tilde{e} = 0.8$). The non-grey regions correspond to the regions of the parameter space for which a SLC is preferable to the baseline (and full commitment).

κ are implementation costs, and thus paid in the last period and discounted accordingly. As a result a lower δ implicitly decreases the cost of the contract which implies that the maximum cost a SLC can incur and still be preferable increases. On the other hand, the cost for the full commitment contract gets always paid in the first period and does not depend on variations in the discount factor. However, while costs remain constant, the relative utility of cooperation induced in the last period decreases with lower discounts factor, which results in a lower maximal cost for the full commitment to be preferable to the baseline.

Lastly, a comment should be made about the cost structure used in the model. It is obvious that focusing solely on implementation costs that do not depend on the length of the contract, its complexity, monitoring or numbers of action prescribed is a stark simplification. However, when comparing the SLC to full commitment contracts, these simplifications not only make the model more tractable. As mentioned before, every aspect mentioned above that does not factor into the height of the costs would make the

depending on $\pi^L \gtrless 0.5$.

cost structure biased in favour of shorter, less restrictive contracts such as the SLC. Thus, the results are based as much as possible on the incentives of cooperation itself rather then cost advantages. Therefore, any generalization of the cost structure, except for the timing of costs, which is examined in Section 5.2 and does not result in any qualitative changes, would merely strengthen the preferability of SLCs over the full commitment contract further.

5 Extensions

In this section, we relax three assumptions from our main analysis. First, we allow for multi-period soft-landing contracts. Second, we consider more general contracting cost structures. And third, we allow players to enter the soft-landing contract at a later stage (rather than the outset of their relationship).

5.1 Multi-Period Soft-Landing Contracts

When considering both effective classes of SLCs (Section 3) before, the analyses were limited to the case where the contract took effect in the last period of the game. This is the most natural case for mainly two reasons. Firstly, compared to the cooperative equilibrium in the baseline, by inducing *intensive* cooperation a SLC can only improve the payoffs in the last period. Since this intensification only takes place in the last period, it appears natural to use the SLC to commit only in the last period as well.[16] Secondly, the SLC comes with costs κ that are paid when the contract takes effect, reflecting implementation costs. Due to discounting, this costs have a present value of $\delta^\tau \kappa$ depending on the period τ in which the contract becomes effective. As a result, having a contract for more than one period is in addition wasteful, as the same cost has to be paid earlier and is thus less discounted.

Even though it appears as if multi-period SLCs can be disregarded on this basis, there is a reason to analyze them; the extensive margin of inducing cooperation. While the intensive margin is concerned with the increase of existing cooperation, the extensive margin covers the creation of new cooperation, i.e., changing the equilibrium from a non-cooperative to a cooperative one.

Consider a SLC as described in Section 3 where the prescription is in effect for a total of τ periods. Both for history-dependent and -independent contracts there is no incentive to deviate once the contract takes effect. On the cost side, the threshold for history-dependent and history-independent SLCs decreases from $\kappa < 1 - \pi^L$ to $\kappa < \delta^{\tau-1}(1 - \pi^L)$, which makes the SLC not only less likely to be implementable but also less favorable than

[16] More specifically, this arises when considering that the commitment of the SLC is used solely to counter the backwards unravelling in the equilibrium, which is restricted to uniquely the last period.

	Cooperative baseline	Non-cooperative baseline
Single-period contract	$\kappa \leq \frac{\delta^T}{1-\alpha+\alpha^T}\left(1-\pi^L\right)$	$\kappa \leq \frac{1-\delta^{T+1}}{(1-\delta)(1-\alpha+\alpha^T)}\left(1-\pi^L\right)$
Multi-period contract	$\kappa \leq \frac{\delta^T}{1-\alpha+\alpha^{T-\tau}}\left(1-\pi^L\right)$	$\kappa \leq \frac{1-\delta^{T+1}}{(1-\delta)\left(1-\alpha+\alpha^{T-\tau}\right)}\left(1-\pi^L\right)$
Full commitment	$\kappa \leq 1-\pi^L$	$\kappa \leq \frac{1-\delta^{T+1}}{1-\delta}\left(1-\pi^L\right)$

Table 2: Cost threshold of the different model specifications for both only implementation costs (main versions) and generalized costs (extension)

the one period version. On the other hand a longer SLC also changes the discount factor necessary for the cooperative equilibrium to be possible. Denote by $\bar{\delta}^i(\tau)$ the respective threshold for the extended SLC.[17] Then, one can show that $\bar{\delta}^i(\tau)$ is weakly decreasing in τ.[18]

Together, this creates a trade-off between the intensive and the extensive margin when extending a SLC. With a longer contractual period the contract becomes implicitly more costly as its present value increases due to less discounting, which reduces the possibility of intensifying cooperation. On the other hand a longer SLC reduces the necessary discount factor for cooperation to be possible in the first place, making extensive cooperation easier to induce. This is possible until the increase in τ lowers $\delta^i(\tau)$ to the point where each player prefers to deviate immediately (even if enduring punishment for the rest of the game), at which point increasing τ is no longer effective and thus never optimal. Thus, for a long enough game, this imposes a natural upper bound on the length of a soft-landing contract - which generically does not last for the entire game.

5.2 Generalized Contract Costs

Up to this point we considered the contractual costs κ to be implementation costs in the sense that they were paid when the contract takes effect. There is, however, a second main type of contractual costs that were ignored so far, *drafting costs*. Drafting costs are costs that emerge from negotiating and setting up a contract and are usually paid ex-ante. To capture both aspects, consider a SLC exhibiting generically both drafting costs κ_d and implementation costs κ_i. To facilitate analysis, let α denote the proportion of total costs taken up by implementation costs, i.e., $\kappa_i = \alpha\kappa$ and $\kappa_d = (1-\alpha)\kappa$.

Using this, we can then reformulate the cost thresholds for the history-independent and history-dependent SLC when compared to both a cooperation and non-cooperation baseline equilibrium. The results are summarized in Table 2.

[17] Under this notation, for instance, $\bar{\delta}^i = \bar{\delta}^i(1)$ for $i = \{IA, DA\}$.

[18] More specifically, it is strictly decreasing as long as $\delta^i(\cdot) > \bar{\delta}^D := \frac{\pi^D-1}{\pi^D}$ and is constant at $\bar{\delta}^D$ afterwards. For a proof of this claim see Appendix A.3.1

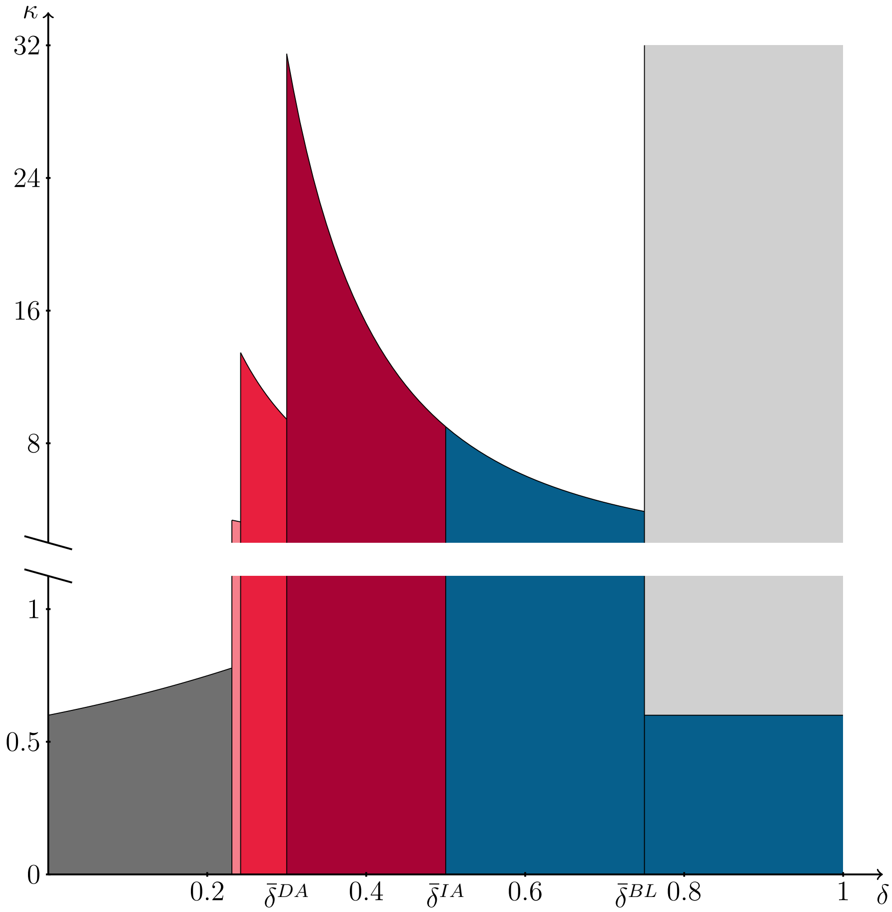

Figure 4: Environments with the highest payoff depending on discount δ and cost κ for the example. As before, in the red area, only the dependent action SLC is implementable, whereas in the blue area both kinds of SLCs are possible. The lighter shades of red represent the additional SLCs with a runtime of 2 and 3 periods.

Figure 5 depicts an overview over the optimality of contracts under two cost regimes (where implementation costs make up half and three-quarters of total costs resp.). As mentioned in Section 3, it can be seen that changing the cost structure to more realistically reflect contracting costs in reality does not change the findings of our main model qualitatively.

The comparison with Figure 3 is useful in order to understand the main difference between different cost structures, namely how optimality (i.e., costs such that a contract is the most preferred option) depends on δ. Without a generalized cost structure, full commitment contracts behave as if they only exhibit drafting costs (as all costs are paid upfront when the contract becomes binding in the first period). In contrast, SLCs only feature implementation costs since no costs are paid upfront. In the former case, a higher discount factor δ does not increase the present value of costs but increases the present value of the payoffs, hence making higher costs feasible. In the latter case, if the contract changes outcomes in any period but the last, those changes are affected less than the changes to the present value in costs, leading to a decline of feasible costs with δ. In

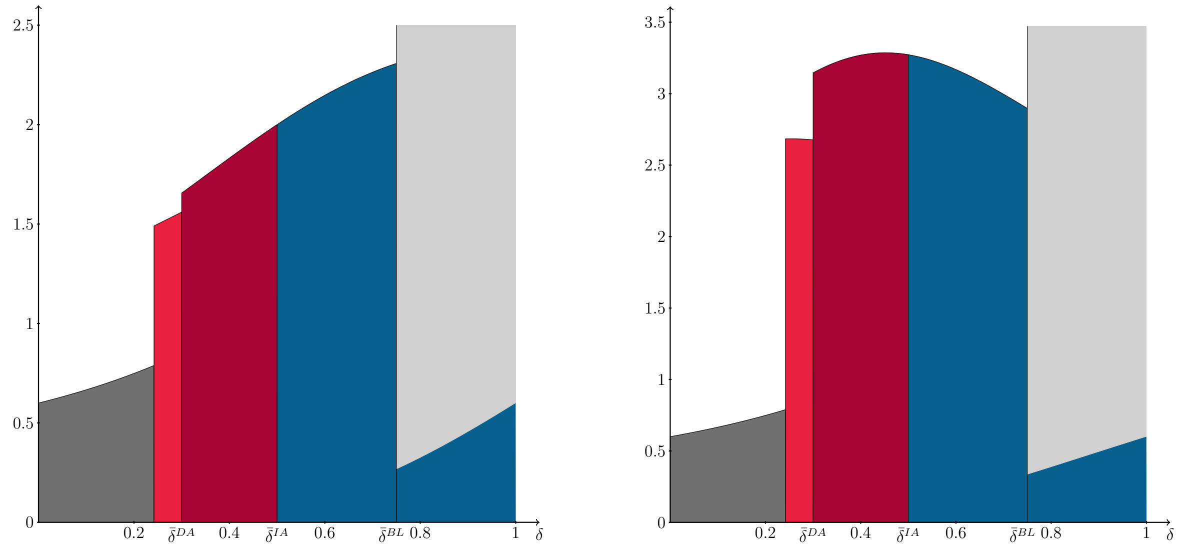

(a) Generalized Costs ($\alpha = 0.5$)

(b) Generalized Costs ($\alpha = 0.75$)

Figure 5: Extensions of Figure 3 including two-period SLCs (in the respective light shade) and drafting costs with $\alpha = 0.5$ (left) and $\alpha = 0.75$ (right).

Figure 5 it becomes visible how generalized costs are affected from changes in δ depending on which effect dominates.

5.3 Endogenous Contracting Timing

When thinking about SLCs in practice there are instances where the contract merely not being effective in every period does not fully capture the contract's nature. Rather, the players start a relationship (i.e., start playing the game in our model) and could postpone negotiating a contract to an eventually later point in time. In principle, this adds an uncertainty to the initial actions that could counteract the incentives of cooperation of any (certain) SLC.

In our main model, contracts have to be formed exogenously in the very first period before any interaction takes place. This section aims to show that this is mainly an assumption that simplifies solving the model with the timing being robust when allowing parties to also form contracts at a later point in time. To capture this, extend the model such that in every period the players, in addition to choosing their action a_i for that period, also decide beforehand whether to enter into any contract.[19] The contract can in principle be of any form highlighted in Section 3 - although we will restrict attention to dependent-action SLCs for expositional ease.

First, it is important to note, that any sequentially rational contracts can only affect payoffs in periods after the contract was negotiated. [20]

[19] As before, we maintain the assumption that contracts negotiated in any stage are symmetric.

[20] Without sequential rationality, contracts could be chosen that are contingent on earlier stages of the game. But crucially once a history is reached in which a contract would need to be punishing, this

For our analysis, two cases need to be distinguished. On one hand, consider the case where $\delta < \bar{\delta}^{BL}$. Following Lemmas 1 & 2 this implies that only the non-cooperative equilibria can exist in the baseline. Importantly, NE are played in every period, making unilateral deviations unprofitable by design. Hence, under the restrictions already discussed in Section 4 whenever a SLC is optimal, writing it as early as possible to extend the stream of cooperation payoffs for as long as possible is optimal.[21]

On the other hand, consider the remaining case where $\delta \geq \bar{\delta}^{BL}$ and start by considering the last period. As stated in Lemma 1 this guarantees the existence of a cooperative baseline, meaning that any SLC can only improve cooperation intensively, i.e., increasing payoffs from π^L to 1 in the last period. Here, in every period but the last players might find it profitable to unilaterally deviate. However, note that SLCs are only optimal as long as $\kappa < 1 - \pi^L$ (see Prop. 3). Under such costs, players will find it optimal to write a contract if no one exists yet both after deviation and non-deviation histories.[22] Since this contract would not condition on past play, it would act similar to the abridged full commitment contract discussed before, practically partitioning the game at the point of the contract, making it optimal to play l in the last period before the contract will be negotiated, thus lowering the total payoffs. The best action in the preceding period is, thus, to write a contract to avoid piL from being reached in this period, effectively moving the partitioning forward until a contract is formed in the very first period, yielding the cooperation payoff in every subsequent period. Put together with the first case, this results in the following result:

Corollary 2 (Endogenous Timing). *Every contract optimal in the basic model with exogenous timing will be also formed in* $t = 0$ *under endogenous timing.*

6 Conclusion

We created a simple framework of symmetric player interaction with potentially partial but not complete cooperation. Introducing contracts into this framework allowed players to extend the cooperation from the baseline both intensively and extensively. While complete contracts (full commitment contracts in our terminology) are effective at doing that, we found that by using soft-landing contracts the same results can be achieved while saving costs.

The main insight consisted in the finding that in order to counteract the backwards unravelling that prevents full cooperation in the absence of contracts, the most cost

contract can not be optimal and is clearly dominated by, e.g. a full commitment contract that treats past play as sunk.

[21] This result extends to the case where costs are of the more general form presented in Section 5.2. For more details, see App. A.3.2.

[22] *Deviation histories* refer to histories of past play in which at least one player deviated at least once from the cooperative action w, with *non-deviation histories* being defined accordingly.

effective contracts are those that become only effective when strictly required to. However, being only in effect for the later stages of the game, SLCs have to be carefully constructed so to not maintain the incentives that led to cooperation in the baseline. We found that abridged full commitment contracts are not able to sustain those incentives whereas SLCs that are either less rigid but incontigent (history-independent SLCs) or contigent (history-dependent SLCs) do. These results proved robust to extensions regarding the cost structure of contracts, the endogeneizing the contract timing as well as relaxing the prescription of actions in favor of specifying (budget-balanced) transfers. Further, we showed how extending the duration of SLCs can be used to trade off cost-effectiveness with implementability, thus inducing even more cooperation in the extensive margin.

Our results contribute to the existing literature on contracts in two major ways. On one hand, they provide a solid justification for the use of such framework contracts that extends beyond the perspective of building and fostering trust, that is usually the focus of the literature on contracting in practice. On the other hand, temporal incompleteness is generally overlooked in the literature on rationally incomplete contracts. While many papers have studied the role of specificity and contingency in contracts and how incompleteness in these dimensions might arise as a conscious decision born out of strategical motivation, this paper is the first, to the best of our knowledge, that undertakes a comparable approach with respect to the temporal dimension.

While this model can serve as a simple starting point to analyze this specific contracting practice, further research is needed to extend the understanding of temporally incomplete, such as developing a framework that generates contracts that are written in ongoing stages of the game. Potential additional avenues include models focusing on asymmetric principal-agent interactions as well as a unified model incorporating the role of trust in negotiating contracts.

References

Anderlini, Luca and Leonardo Felli (1994) "Incomplete Written Contracts: Undescribable States of Nature," *Quarterly Journal of Economics*, 109 (4), 1084–1124.

Arrighetti, A., R. Bachmann, and S. Deakin (1997) "Contract Law, Social Norms and Inter-Firm Cooperation," *Cambridge Journal of Economics*, 21 (2), 171–195.

Axelrod, Robert (1981) "The Emergence of Cooperation Among Egoists," *American Political Science Review*, 75, 306–318.

Basov, Suren (2016) *Social Norms, Bounded Rationality, and Optimal Contracts*, 30 of Studies in Economic Theory, Singapore: Springer.

Battigalli, Pierpaolo and Giovanni Maggi (2002) "Rigidity, Discretion, and the Costs of Writing Contracts," *American Economic Review*, 92 (4), 798–817.

——— (2008) "Costly Contracting in a Long-Term Relationship," *RAND Journal of Economics*, 39 (2), 352–377.

Ben-Ner, A. and L. Putterman (2009) "Trust, Communication and Contracts: An Experiment," *Journal of Economic Behavior & Organization*, 70 (1), 106–121.

Benoit, Jean-Pierre and Vijay Krishna (1985) "Finitely Repeated Games," *Econometrica*, 53 (4), 905–922.

Burchell, Brendan and Frank Wilkinson (1997) "Trust, Business Relationships and the Contractual Environment," *Cambridge Journal of Economics*, 21 (2), 217–237.

Diamond, Douglas W. (1984) "Financial Intermediation and Delegated Monitoring," *Review of Economic Studies*, 51 (3), 393–414.

——— (1991) "Monitoring and Reputation: The Choice between Bank Loans and Directly Placed Debt," *Journal of Political Economy*, 99 (4), 689–721.

Dye, Ronald A. (1985) "Costly Contract Contingencies," *International Economic Review*, 26 (1), 233–250.

Fudenberg, Drew and Jean Tirole (1991) *Game Theory*: MIT press.

Holmstrom, Bengt and Paul Milgrom (1991) "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law, Economics, and Organization*, 7 (Special Issue), 24–52.

Kaplow, Louis (1995) "A Model of the Optimal Complexity of Legal Rules," *Journal of Law, Economics, and Organization*, 11 (1), 150–163.

Klein, Benjamin (2002) "The Role of Incomplete Contracts in Self-Enforcing Relationships," in Brousseau, Eric and Jean-Michel Glachant eds. *The Economics of Contracts: Theories and Applications*, 59–71, Cambridge, UK: Cambridge University Press.

Kreps, David M., Paul Milgrom, John Roberts, and Robert Wilson (1982) "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27 (2), 245–252.

Lorenz, E. (1999) "Trust, Contract and Economic Cooperation," *Cambridge Journal of Economics*, 23 (3), 301–315.

Malhotra, Deepak and J. Keith Murnighan (2002) "The Effects of Contracts on Interpersonal Trust," *Administrative Science Quarterly*, 47 (3), 534–559.

Melumad, Nahum, Dilip Mookherjee, and Stefan Reichelstein (1997) "Contract Complexity, Incentives, and the Value of Delegation," *Journal of Economics and Management Strategy*, 6 (2), 257–289.

Ohmae, Kenichi (1989) "The Global Logic of Strategic Alliances," *Harvard Business Review*, March-April, 143–154.

Schwartz, Alan and Joel Watson (2004) "The Law and Economics of Costly Contracting," *Journal of Law, Economics, and Organization*, 20 (1), 2–31.

Tippins, Karl (2021) "Steps to Building Trust with Contracts," Published online at https://contractbook.com/blog/steps-to-building-trust-with-contracts (November 22, 2021).

Tirole, Jean (2009) "Cognition and Incomplete Contracts," *American Economic Review*, 99 (1), 265–294.

Williamson, Oliver E. (1985) *The Economic Institutions of Capitalism*, New York, NY: Free Press.

——— (1989) "Transaction Cost Economics," in Schmalensee, Richard and Robert Willig eds. *Handbook of Industrial Organization*, 1, Chap. 3, 135–182, Amsterdam, NL: North-Holland.

A Appendix

A.1 Proofs

Proof of Lemma 1

Proof. Consider first the timing of the optimal deviation. To that end, assume w has been played for the first τ periods and consider the indifference to deviate in $t = \tau + 1$ compared to $t = \tau + 2$, yielding

$$\begin{aligned} &\frac{1-\delta^\tau}{1-\delta} + \delta^\tau \pi^D = \frac{1-\delta^\tau}{1-\delta} + \delta^\tau + \delta^{\tau+1}\pi^D \\ \iff &\delta^\tau \pi^D = \delta^\tau + \delta^{\tau+1}\pi^D \\ \iff &\pi^D = 1 + \delta\pi^D \iff \delta = 1 - \frac{1}{\pi^D} = \frac{\pi^D - 1}{\pi^D} =: \bar{\delta}^D \end{aligned}$$

By induction, logic follows that whenever $\delta < \bar{\delta}^D$, it is optimal to deviate in $t = 1$, whereas for $\delta > \bar{\delta}^D$ deviation in $t = T$ is the most profitable.

For $\delta < \bar{\delta}^D$, the condition for the existence of a cooperative equilibrium gets more complicated as indifference is characterized by

$$\frac{1-\delta^T}{1-\delta} + \delta^T \pi^L = \pi^D$$

However, the following argument shows that, in this case, no cooperative equilibrium can exist. Assume that a player gets a utility stream of $\frac{1}{1-\delta}$, which is equivalent to the infinite sum of discounted payoffs of $\pi_t = 1$, i.e., $\sum_0^\infty \delta^t = \frac{1}{1-\delta}$. It is easy to see that getting $\pi_t = 1$ for infinite periods has to yield a higher total utility than getting $\pi_t = 1$ for only $T-1$ periods and $\pi_t = \pi^L < 1$ in the last. Note, however, that this infinite stream is equivalent to the first-period deviation payoff whenever

$$\frac{1}{1-\delta} = \pi^D \iff \pi^D(1-\delta) = 1 \iff \delta = \frac{\pi^D - 1}{\pi^D} =: \bar{\delta}^D$$

Since the infinite stream utility increases with δ, it has to be that $\pi^D \geq \dfrac{1}{1-\delta} > \dfrac{1-\delta^T}{1-\delta} + \delta^T \pi^L$ whenever $\delta < \bar{\delta}^D$. But that is saying that deviation is always profitable compared to cooperation. Thus, the only cooperative equilibria can arise when $\delta \geq \dfrac{\pi^D - 1}{\pi^L} > \dfrac{\pi^D - 1}{\pi^D}$ since $\pi^D > 1 > \pi^L$. □

A.2 Mixed Equilibria of the APD

Proof. To check for mixed equilibria of the stage game, first note that w is dominated by s by assumption:

1) $0 > 2(\pi - c) = 2(\frac{1}{2} - \pi)$ by (a)
2) $2\pi - m > 4\pi - 2c = 1 \iff m < 2\pi - 1by$ by (c)
3) $\tilde{e}\pi - m > (2+\tilde{e})\pi - 2c = (2+\tilde{e})\pi - 4\pi + 1 \iff m < 2\pi - 1$ by (c)

Hence the only mixed equilibrium is given by the mixed strategy $\alpha_i = ps + (1-p)l$. The indifference condition for mixed equilibria then yields $p = \frac{\tilde{e}}{m}(\pi - \frac{1}{2})$. Since $\pi_i(s,l), \pi_i(l,s) < 1/2\tilde{e}$, it follows that the utility of the mixed NE is below π^L.[23] Finally, the alternative assumption $(d)'$ is obtained from the condition ensuring $\tilde{e}\pi - m > 0$, which in turn implies a positive payoff from the mixed NE.

Without this assumption, the MNE could provide a negative payoff and thus offer an alternative to punishing deviation with playing (s, s). This would affect the δ-thresholds but not change the result qualitatively. □

A.3 Formal Derivations of Extensions

A.3.1 Multi-Period SLC

To show that $\bar{\delta}^i(\tau)$ is decreasing in τ, assume for now a deviation to be always optimal in the last period, and consider first the case of a history-independent SLC lasting for τ periods, $\bar{\delta}^{IA}(\tau)$. Implicitly, cooperating is optimal if

$$1 + \delta + \ldots + \delta^\tau \geq \pi^D + (\delta + \ldots + \delta^\tau)\pi^L$$
$$\iff \quad \delta + \ldots + \delta^\tau \geq \frac{\pi^D - 1}{1 - \pi^L}$$

which is strictly increasing in δ on the LHS and constant on the RHS. Thus, for this to hold, we need $\bar{\delta}^i(\tau)$ to be decreasing in τ.

Now consider the feasibility threshold for a history-dependent SLC of τ periods, $\bar{\delta}^{DA}(\tau)$, which is implicitly given by

$$\pi^D - 1 < \delta + \ldots + \delta^\tau$$

and note that for $\tau + 1$ periods the threshold is characterized by

$$\pi^D - 1 < \delta + \ldots + \delta^\tau + \delta^{\tau+1} \iff \pi^D - 1 - \delta^{\tau+1} < \delta + \ldots + \delta^\tau$$

Since $\delta^{\tau+1} > 0$ for all $\delta > 0$ and the RHS is strictly increasing in δ for $\delta > 0$, the only possibility for the inequality to hold is to have δ decrease, which is saying $\bar{\delta}^{DA}(\tau) > \bar{\delta}^{DA}(\tau + 1)$. By induction, this implies that $\bar{\delta}^{DA}(\tau)$ is strictly decreasing in τ.

Lastly, we now relax the assumption that deviation is always optimal in the last period. As seen in Appendix A.1, an early deviation is (strictly) optimal whenever $\delta < \bar{\delta}^D := \frac{\pi^D - 1}{\pi^D}$ which implies

$$\pi^D > 1 + \ldots + \delta^T - 1 + \delta^T \pi^D > 1 + \ldots \delta^T$$

Hence, no contract that gives the optimal (w, w)-payoff in each period can be optimal and thus $\bar{\delta}^i(\tau) \geq \bar{\delta}^D$ for all τ.

For the case of implementation costs in the non-cooperative baseline setting for $\tau = n^c$ periods, we obtain

$$\sum_{t=0}^{T} \delta^t - \delta^\tau \kappa > \pi^L \sum_{t=0}^{T} \delta^t \iff \kappa < \frac{(1 + \ldots + \delta^T)}{\delta^\tau}(1 - \pi^L) = \frac{1 - \delta^{T+1}}{(1 - \delta)\delta^\tau}(1 - \pi^L)$$

for the SLC to be preferable.

[23] For $\pi_i(s, l)$ this is satisfied since (l, l) is a NE, for $\pi_i(l, s)$ this follows from the fact that $\pi_i(l, s) = \tilde{e}(\pi - c) < \tilde{e}(2\pi - c) = \frac{1}{2}\tilde{e}$ since $\tilde{e}, \pi > 0$.

A.3.2 Endogenous Timing with Generalized Costs

As in Section 5.2, assume that costs are a combination of drafting and implementation costs with $\kappa_i = \alpha\kappa$.
Consider first the case in which only the non-cooperative baseline exists ($\delta < \bar{\delta}^{BL}$). Next, consider a SLC becoming active in period τ and compare it to one that becomes active in period $\tau + 1$. For the player to prefer the former, we require

$$\begin{aligned}
&\delta^{\tau-1}(1-\kappa_d) + \delta^{\tau} > \delta^{\tau-1}\pi^L + \delta^{\tau}(1-\kappa_d) \\
&\iff 1 - \kappa_d + \delta > \pi^L + \delta(1-\kappa_d) \iff 1 - \pi^L > \kappa_d(1-\delta) \iff \kappa_d < \frac{1-\pi^L}{1-\delta}
\end{aligned}$$

which is independent of τ and, hence, by induction any SLC is either optimal in $t = 0$ or $t = T$. Now consider a SLC optimal in the last stage and compare it to the non-cooperative baseline. The latter is optimal whenever $\delta^T\pi^L > \delta^T(1-\kappa)$ which is equivalent to $\kappa > 1 - \pi^L$ which is necessarily true since $\kappa > \kappa_d > \frac{1-\pi^L}{1-\delta} > 1 - \pi^L$. Hence, either a SLC is negotiated at $t = 0$ or no SLC is formed, which simplifies to the standard model of Section 3.

Now consider the case where $\delta \geq \bar{\delta}^{BL}$ implying the existence of a cooperative equilibrium. Here we need to consider the possibility of unilateral deviation in intermediate period, hence approach the game using backward induction. In the last period, following a non-deviation history the payoff without contract is π^L while after a deviation-history it is 0. Under an existing contract, the payoff is $1 - \kappa$.[24] Thus players would prefer to form a contract in all histories if $\kappa < 1 - \pi^L$ and only in deviation histories if $\kappa < 1$.[25] As discussed before, these contracts need to be independent of histories of prior play in order to be sequentially rational. Hence, the last period yields $1 - \kappa$ irrespective of whether players deviated previously. Hence, in period $T - 1$ playing (w, w) is no longer optimal since players would deviate and still enter a non-punishing contract afterwards. Hence players in a non-deviation history face the choice of sticking to the second-best π^L or forming a contract that gives $1 - (1-\delta)k_d$ and in both cases getting $1 - \kappa$ in the last period. For not contracting immediately, we would need $\pi^L > 1 - (1-\delta)\kappa_d$ which is equivalent to $\kappa_d > \frac{1-\pi^L}{1-\delta}$ and contradicting $\kappa < 1-\pi^L$. Hence, it is optimal to contract. It follows that the same must hold for players in a deviation history facing a zero payoff without contract. Hence again, after both kinds of histories a contract would be made partitioning the game in the same way as before. Crucially know, in each preceding history, the highest attainable payoff without contract is π^L after non-deviation histories, 0 after deviation histories. The last step is to prove that in any period a contract adds $1 - (1-\delta)\kappa_d$. This will be proven by complete induction. Start from a last period contract in which the total payoff is $1-\kappa$ trivially. Now take an arbitrary contract that yields $1 - \kappa_d + \delta + \ldots + \delta^{n-1} + \delta^n(1-\kappa_i)$ and note that we obtain the previous period total value of $1 - \kappa_d + \delta + \ldots + \delta^n + \delta^{n+1}(1-\kappa_i)$ by adding $1 - (1-\delta)\kappa_d$ to the discounted value, i.e.,

$$\begin{aligned}
&\ 1 - (1-\delta)\kappa_d + \delta(1 - \kappa_d + \delta + \ldots + \delta^{n-1} + \delta^n(1-\kappa_i)) \\
=&\ 1 - \kappa_d + \delta\kappa_d + \delta - \delta\kappa_d + \ldots + \delta^n + \delta^{n+1}(1-\kappa_i) \\
=&\ 1 - \kappa_d + \delta + \ldots + \delta^n + \delta^{n+1}(1-\kappa_i)
\end{aligned}$$

Thus, the game transforms for each subsequent period into a stationary where players decide between π^L and $1 - (1-\delta)\kappa_d$ in non-deviation histories and 0 and $1 - (1-\delta)\kappa_d$ in deviation

[24] Technically, it is only $1 - \kappa_i$ if the contract was drafted in earlier periods, but this assumption simplifies analysis and will be accounted for in later stages.

[25] Since we have shown in Prop. 3 that $\kappa < 1 - \pi^L$ is required for SLCs to be optimal at all, we can rule out an SLC becoming profitable if it isn't in the last period (the reasoning propagates forward until the initial node is a unique non-deviation history at which not contracting is optimal). Hence, focus on the former case where $\kappa < 1 - \pi^L$.

histories, followed by a safe contract irrespective of their choice in that period. Thus, the only SLCs that arise are when $\kappa < 1 - \pi^L$ and are agreed upon in $t = 0$.

A.4 Transfer-Based SLC

Another approach to model SLCs in this framework is to think of a contract of a commitment to transfer payoff between players rather than directly modifying the action space. In the model setting, this means that a SLC corresponds to a change of the payoff of a stage game. For simplicity, we will focus again solely on symmetrical contracts. For now, consider transfers that are not conditioned on previous actions. Then for one period t a history-independent transfer SLC can be defined as

$$C_t : A_i \times A_i \to \mathbb{R} \text{ such that } \forall a_i, a_j : C_t(a_i, a_j) + C_t(a_j, a_i) \leq 0.$$

resulting in the total utility of an action profile, $u_{i,t}(a_{i,t}, a_{j,t}) = \pi_i(a_{i,t}, a_{j,t}) + C_t(a_{i,t}, a_{j,t})$. The budget condition ensures that transfers can not create additional surplus (*feasibility*).[26] Under feasibility, transfer SLCs can replicate any action prescription SLC by simply setting $C_t(a_i, \cdot) < -T\pi^D$ for any action profile a_i that should not be played, making the transfer costlier than the maximum payoff obtainable in the entire game. We will focus therefore on the subset of budget-balanced transfers, where $C_t(a_i, a_j) = -C_t(a_j, a_i)$, with the additional advantage of not wasting social surplus.

Given that $\pi_i(w, w) = 1$ is only exceeded by $\pi^D > 1$, (w, s) and (s, w) are the sole candidates for improving upon (w, w). However, $\pi_i(w, s) + \pi_j(w, s) < 2$, which makes (w, w) the socially optimal action profile, even when considering all feasible transfers.
Therefore, focus on transfers that can implement (w, w) in each period. Note that a credible last period strategy has to induce a NE of the stage game. Therefore, the contract needs to make (w, w) a NE of the stage game in (at least) the last period. This is easily achieved by setting $C_t(w, l) = 1 - \pi^D - \varepsilon$ where $\varepsilon > 0$. In order to keep (s, s) a NE to use in the grim-trigger strategy, ε has to be sufficiently small.
In order to prove the existence of an ε that induces (w, w) in each period, first consider optimality of (w, w) and note that

$$\pi_i(w, s) + \pi_j(w, s) = 2\pi - m + 2(\pi - c) = 4\pi - 2c - m = \pi_i(w, w) - m < 2\pi_i(w, w) = 2$$

with $\pi_i(w, w) = 4\pi - 2c = 1$ by design.

Then, a contract that induces (w, w) as an equilibrium needs to fulfill two conditions, Firstly, (w, w) has to become a NE of the stage game, requiring

$$\pi_i(w, s) + C_t(w, s) < 1 \iff C_t(w, s) < 1 - \pi^D \iff C_t(w, s) = 1 - \pi^D - \varepsilon \text{ with } \varepsilon > 0$$

Secondly, (s, s) has to remain a NE of the stage game. Consider again the inequality above and note that

$$\begin{aligned} \pi_i(s, w) + C_t(s, w) &= \pi_j(w, s) - C_t(w, s) = \pi_j(w, s) + \pi_i(w, s) + \varepsilon - 1 \\ &= \pi_i(w, w) - 1 + \varepsilon - m = \varepsilon - m \end{aligned}$$

Thus, $\pi_i(s, w) + C_t(s, w) < 0$ as long as $\varepsilon < m$. Since $m > 0$ by definition, this implies that there always exists an ε satisfying $0 < \varepsilon < m$.

[26]Unless $C_t(a_i, a_j) + C_t(a_j, a_i) = 0$ (*budget-balance*) it is, however, possible to destroy surplus.

Intuitionally, if ε is too high, that is, the transfer paid by the shirking to the working player too big, it deters not only free-riding, but also incentivizes players to work when the opponent is shirking. Thus, the simplest budget-balanced transfer SLC that implements (w, w) as a stage NE is given by

$$C_t(a_i, a_j) = \begin{cases} 1 - \pi^D - \varepsilon & \text{for } (w, s) \\ \pi^D - 1 + \varepsilon & \text{for } (s, w) \\ 0 & \text{otherwise} \end{cases}$$